5-29-2024 10:00 AM

Essays on Firms and Climate Change

Emmanuel Murray Leclair, *Western University*

Supervisor: David Rivers, *The University of Western Ontario*
: Salvador Navarro, *The University of Western Ontario*
: Sergio Ocampo Diaz, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Economics

Abstract

This thesis comprises three chapters examining firms' and markets' responses to regulation within the context of climate change.

In Chapter 2, I analyze how steel firms in India would respond to counterfactual carbon pricing, revealing a 70% emissions reduction for a carbon tax equivalent to 2,000 INR/ton (25 USD/ton) of carbon dioxide equivalent (CO_{2e}). By developing and estimating a rich production model, I find that only 18% of this reduction stems from fuel-switching within firms, while the majority comes from output reallocation across firms. This is because firms are differentially exposed to the tax, which increases the competitiveness of cleaner firms at the expense of dirtier firms. Chapter 3 investigates how firms adjust production in response to asymmetric carbon pricing across Canadian provinces. Using Canadian plant-level pollution release data, I find evidence that carbon leakage mitigated 45% of emissions reduction efforts in British Columbia and Quebec, following two carbon taxes implemented in 2008 and 2007, respectively. Regulated firms became less competitive, and output shifted to unregulated firms in other Canadian provinces. In Chapter 4, I assess the efficiency of the U.S. natural gas pipeline network in responding to climate-induced demand and supply shocks. Analysis of the February 2021 cold wave in Texas reveals that an unregulated secondary market effectively reallocates capacity to alleviate supply and demand imbalances. However, concerns arise regarding market power on both sides of the market, potentially reducing efficiency gains from this secondary market.

Keywords Climate change, firm dynamics, input choice, fuel efficiency, production function, carbon leakage, pipeline networks

Summary for Lay Audience

This thesis explores the interplay between firms and climate change. Specifically, I study how firms can reduce greenhouse gas emissions in response to environmental policy and investigate some of the unintended consequences of these policies. I also study the consequences of climate change on energy markets.

In Chapter 2, I study how firms respond to carbon pricing in a context where firms with very different technologies make decisions about which fuels to use and how much of each fuel to burn. I build a novel model to recover these fuel choices and quantify the impact of carbon pricing in the Indian steel industry using detailed plant-level data. I find that implementing a carbon tax equivalent to 2,000 INR/ton (25 USD/ton) of carbon dioxide leads to a 70% reduction in emissions. But only 18% of this reduction comes from fuel-switching within existing firms. I find that the larger reductions come from the reallocation of output across firms (58%) and the reduction in aggregate output (24%).

While firms compete with each other across geographical regions, carbon policy is not always uniform across regions. This can lead to carbon leakage, shifting emissions from regulated to unregulated regions. In Chapter 3, I build a model that allows for region-specific carbon taxes. I estimate the model with publicly available data on a wide range of pollutants to quantify the effect of the British-Columbia (B.C.) carbon tax implemented in 2008. I find strong evidence of carbon leakage in other Canadian provinces, mitigating 45% of emissions reduction efforts in B.C. and Quebec.

In Chapter 4, I study the efficiency of the U.S. natural gas pipeline network in allocating pipeline capacity in response to unexpected weather shocks, increasingly common with climate change. I find that an unregulated secondary market, where contract holders can lease capacity to other shippers, reacts to significant regional demand fluctuations induced by these weather events, and alleviates the imbalance between supply and demand. However, a largely unregulated secondary market within a heavily regulated primary market raises concerns about market power.

Co-Authorship Statement

Chatper 4 was written alongside Yanyou Chan, Assistant Professor at the University of Toronto, and Adam Wyonzek, Associate at The Brattle Group. This project, is a collaborative effort. From the very beginning and until the present version, each of us contributed equally.

Acknowledgements

I am filled with gratitude to my supervisor, David Rivers, who supported me, accepted my ideas, and guided my research to what it is today. David was always willing to meet with me and help me see through what would often feel like a web of complex details. When faced with the unsurmountable—or at least what I thought—he would often tell me to keep trying and to avoid reallocating my time to easier tasks. To my surprise, it worked. His relentless eargerness to make sure that I keep my focus, one paper at a time, was just what I needed. I hope to be such a mentor to my students one day.

I am also greatly indebted to my committee members, Salvador Navarro and Sergio Ocampo Diaz. Both are gifted with incredible wit and capable of understanding my work better than myself within a few minutes of explanation. Their feedback greatly contributed to the methodological depth of my work and my ability to share it with others.

This section would not be complete without my appreciation for the faculty at Western. To my surprise, it seems to be that the entire department was enthusiastic about my work and provided me with their time and the opportunity to present my work again and again. Despite the countless times I presented, I was faced with renewed feedback every time. I want to extend a special thanks to Nirav Mehta, Tim Conley, Rory McGee, Baxter Robinson, Lance Lochner, Daniel Chaves, Ananth Ramanarayanan, Todd Stinebrickner, Lisa Tarquinio, Nail Kashaev, and Michael Sullivan for their feedback. I also want to thank Charles Sauders, who passed away recently. Charles read my paper many times, and provided me with feedback unique to himself. He was always keen to help and will be missed.

I have made countless professional relationships along the way and have been fortunate to attend many conferences in Canada, the U.S., and Europe. I want to extend a particular thanks to Yanyou Chen and Adam Wyonkez, who have been and are still superb coauthors.

I also want to express my gratitude to my friends and colleagues at Western. As PhD students, we would often exchange ideas, but most importantly, we would hang out and escape, if only for a while, the all-consuming effect that research has on us. So thank you Tommas Trivieri, Anastasiia Suvorova, Evan Sauve, Francisco Martinez Hernandez, Karthik Kotikalapudi, Martin Bustos, Matt Carew, Giao Dang, Emily Hanson, Juan Holguin, Jihyun Kim, Nalinda Murray, Henk Schouten, Martin Luccioni, Robert Millard, Marco Pariguana, and Ali Kamranzadeh. My time in London (Ontario) wouldn't have been the same without you. You have contributed to some of the best memories of my life.

Lastly, I want to thank my mother, Chantal Murray, my father, André Leclair, and my brother, Louis Murray Leclair. Despite the distance, they have been the keenest supporters of my life choices and make me proud.

Table of Contents

List of Tables

List of Figures

List of Appendices

Chapter 1

Introduction

This thesis explores the interplay between firms and climate change. Specifically, I study how firms can reduce greenhouse gas emissions in response to market-based environmental policy and investigate some of the unintended consequences of these policies. I also study the resilience of energy markets to challenges posed by climate change, such as the increasing advent of extreme weather events.

In Chapter 2, I study firms' heterogeneous and dynamic responses to carbon pricing. The economic cost of carbon pricing depends on the ability and incentives of firms to switch towards cleaner fuels. Yet, many fundamental economic forces that drive firms' decisions to use different fuels are unobserved, causing significant uncertainty over the effectiveness of carbon policies. I propose a new dynamic production model with multidimensional unobserved heterogeneity that underly technology differences across firms and captures how firms' fuel choices respond to price changes. These differences create heterogeneity in abatement costs, which generates heterogeneous responses to carbon pricing. Leveraging minimal assumptions about optimal input choice and the technology frontier, I quantify the model from a detailed panel of Indian steel establishments. Based on these estimates, implementing a carbon tax equivalent to 2,000 INR/ton (25 USD/ton) of CO_{2e} leads to a 70% reduction in emissions. However, only 18% of this reduction comes from fuel-switching within existing firms. I find that the larger reductions come from the reallocation of output across firms (58%) and the reduction in aggregate output (24%).

Motivated by firms's responses to carbon pricing in the previous Chapter, I study how firms reorganize their production when faced with asymmetric carbon pricing in Chapter 3. While firms compete with each other across geographical regions, some regions may put a tax on carbon while others do not. This can lead to carbon leakage, shifting emissions from regulated to unregulated regions. I build a model of imperfect competition with multiple fuels as energy inputs, which allows for region-specific carbon taxes. I estimate the structural model with publicly available Canadian plant-level data on a wide range of pollutants emitted in the air to quantify the effect of the British-Columbia (B.C.) and Quebec carbon taxes implemented in 2008 and 2007, respectively. I find strong evidence of carbon leakage in other Canadian provinces, which mitigated 45% of emissions reduction efforts from the policies in B.C. and Quebec. Similarly to the finding in Chapter 2, I show that Canadian plants did not find it profitable to switch between fuels. As a result, regulated firms became less competitive, and output was reallocated towards unregulated firms.

In Chapter 4, I transition from studying how markets can mitigate climate change to studying challenges posed by climate change and how these challenges affect markets' resilience. Specifically, I assess the efficiency of the U.S. natural gas pipeline network in allocating pipeline capacity in response to unexpected demand and supply shocks, which are increasingly common with climate change. Long-term contracts, the main mechanisms to allocate capacity, can create an imbalance between supply and demand because they reserve capacity for contract holders prior to shocks. Utilizing daily transaction data from 2005 to 2023, I find that an unregulated secondary market, where contract holders can lease capacity to other shippers, reacts to significant regional demand fluctuations and alleviates the imbalance between supply and demand. I also find evidence that forming long-term relationships between buyers and sellers reduces search costs in the secondary market. However, a largely unregulated secondary market within a heavily regulated primary market raises concerns about market power. Evidence shows that buyer and sellers' market shares influence prices, reducing the secondary market's efficiency gains. This secondary market is a rare, real-world example of a market for legal entitlements, and my study is the first to examine it empirically. I provide some evidence that generalizes to similar heavily regulated markets with large sunk investment costs, and I provide guidance for the organization of pipeline networks beyond the U.S.

Chapter 2

Balancing Production and Carbon Emissions with Fuel Substitution

2.1 Introduction

The reliance on fossil fuels in many manufacturing processes has profound environmental repercussions. The release of greenhouse gases from fossil fuel combustion is the largest contributor to climate change, accounting for 75% of total greenhouse gas emissions (UN Climate, 2023). Market-based policies such as carbon taxes and tradable permit systems have been proposed to induce firms to internalize the social cost of fuel combustion. One of the main objectives of these policies is to incentivize firms to adopt cleaner fuels, thereby reducing emissions while maintaining stable levels of production. The economic costs of such policies, such as loss of output and higher consumer prices, thus crucially depend on the ability and incentives of firms to substitute cleaner fuels for dirtier fuels.

Fundamental economic forces that drive firms' decisions to use different fuels are ubiquitous. They reflect disparities in prices, technology, and access to critical infrastructures like transportation networks, all of which vary across space and over time (Collard-Wexler and De Loecker, 2015, Scott, 2021). Yet, little is known about how these forces interact with each other, partly because they reflect unobserved heterogeneity and forward-looking considerations that are difficult to quantify. The incentives of firms to adopt a new fuel in response to a policy may vary substantially based on private information about how efficiently they expect to use this new fuel. Similarly, the ability of firms to substitute between existing fuels may be limited or enabled by specific manufacturing processes.

This paper proposes a new dynamic production model that combines multiple fundamental forces to understand how firms make fuel choices and how these choices respond to price changes. This model differs from existing production models of interfuel substitution in that it captures plants' multidimensional fuel choices, fuel productivity heterogeneity due to unobserved technology differences, and dynamic switching between fuel sets subject to fixed switching costs.[1] Leveraging the approach pioneered by Carlson, Burtraw, Cropper and Palmer (2000) and Atkinson and Luo (2023), I show that these factors provide a microfoundation for abatement costs heterogeneity that reflects the ability of firms to substitute cleaner fuels for dirtier fuels in the short and long run. Combined together, these factors thus generate heterogeneous and dynamic responses to price changes that have important implications for carbon policy.

While there is a significant body of work estimating production functions, understanding the heterogeneity and dynamics underlying firms' input selection when they can choose from multiple input combinations poses new challenges.[2] One of these challenges is quantifying input-specific productivity for an input a firm has never used. Leveraging minimal assumptions about optimal input choices and the technology frontier along with a detailed panel of Indian steel plants between 2009 and 2016, I quantify the rich heterogeneity in plants' incentives and ability to substitute between fuels. The panel features detailed information about plant-specific input prices and quantities along with location within 775 districts, allowing me to interpret this heterogeneity in the context of plants' proximity to key infrastructures such as coal mines and natural gas pipeline networks.

I then use the model to predict how Indian steel plants would respond to various carbon taxes over a horizon of 40 years. India is the second largest steel producer, and steel is one of India's most polluting industries, with coal accounting for nearly 70% of its energy sources. I perform counterfactual policy simulations by imposing a carbon tax levied on fossil fuels, varying the level of the tax. I find that cutting emissions by 50% relative to a no-tax scenario entails a reduction of output by 6.5%. In contrast, under an economy in which all fuels are assumed to be equally productive across firms, as is standard in the production function and abatement costs literature (Atkinson and Luo, 2023, Carlson, Burtraw, Cropper and Palmer, 2000, Fowlie, Reguant and Ryan, 2016, Hawkins-Pierot and Wagner, 2022, Shapiro and Walker, 2018), I find that obtaining the same 50% reduction in emissions leads to a 12% reduction in output, thus overestimating the output costs of such a policy by almost 100%.

[1] The vast majority of papers on interfuel substitution consider firms' fuel sets to be fixed in time and do not account for fuel-specific productivity. See Cho, Nam and Pagan (2004), Ganapati, Shapiro and Walker (2020), Hyland and Haller (2018), Ma, Oxley, Gibson and Kim (2008), Pindyck (1979), Wang and Lin (2017)

[2] See Ackerberg, Caves and Frazer (2015), Blundell and Bond (2000), Demirer (2020), Gandhi, Navarro and Rivers (2020), Grieco, Li and Zhang (2016), Levinsohn and Petrin (2003), Olley and Pakes (1996), Zhang (2019).

Multidimensional heterogeneity significantly reduces the predicted economic cost associated with reducing emissions because it improves the ability of the carbon price to target high-emission plants. Plants that are more productive at dirty fuels relative to cleaner fuels have a higher marginal abatement cost because they are less willing to substitute away from their most productive fuels. Consequently, they face a larger increase in their marginal cost, pass a larger portion of the tax to consumers, and become less competitive relative to plants that are more productive at using cleaner fuels. My results thus provide quantitative support for the long-established idea in environmental economics that firm-level heterogeneity in abatement costs improves the effectiveness of market-based mechanisms relative to command-and-control regulation (Goulder and Parry, 2008).

In the Indian steel context, the carbon tax generates a significant reallocation of output from plants using large coal-powered blast furnaces in Eastern India to cleaner plants in Western India. This reallocation of output is the main channel for reducing emissions. With a carbon price equivalent to $25 U.S. dollars per ton of carbon, 58% of emissions reduction comes from output reallocation, 24% from an aggregate reduction in output, and only 18% from plant-level fuel substitution. Large degrees of fuel specialization due to fuel productivity and high fuel switching costs explain the lack of fuel substitution, capturing significant technology lock-in. To tackle this technological lock-in, I study a subsidy that reduces the fixed cost of natural gas adoption, and I find that it is not a cost-effective tool for reducing emissions. I find that 90% of subsidy beneficiaries are inframarginal plants that would have switched to natural gas regardless, significantly increasing the cost of incentivizing a meaningful change in natural gas adoption. I estimate the cost to increase natural gas adoption by five percentage points at the equivalent of $2.793 billion U.S. dollars over 40 years, amounting to 12% of the industry's profits over the same period.

Quantifying the role of fossil fuels in production in this context requires overcoming two important measurement issues. First, the energy plants use in production is unobserved because it depends on how plants use fuels. The energy service a plant receives differs from the quantity of fuels it uses, measured in common heating potential units. The wedge between a fuel's heating potential and the energy service it provides reflects the fuel's productivity. Second, plants choose fuel sets on the basis of unobserved heterogeneity, such as how productive they would be at using alternative fuel combinations. For instance, a plant can choose to use coal because it anticipates high coal productivity and low gas productivity. However, the plant's gas productivity remains unobserved to the researcher.

I address these measurement issues in three steps. First, I identify the latent quantity and price of energy services by adapting the methods of Ganapati, Shapiro and Walker (2020), Grieco, Li and Zhang (2016). This method relies on optimality conditions from profit maximization to map observed relative input spending to unobserved relative input quantities. Second, I estimate the function that maps fuels to energy services following Zhang (2019) and Blundell and Bond (1998, 2000). This allows me to recover the distribution of fuel productivity across plants. Third, I adapt Arcidiacono and Jones (2003) to jointly estimate fixed switching costs and the distribution of fuel productivity for unused fuels. Using this three-step approach, I recover all production function parameters, the distribution of fuel productivity, and switching costs between fuel sets. Lastly, I estimate the elasticity of substitution between output varieties following Ganapati et al. (2020). These estimates allow me to conduct policy counterfactuals that affect plants' fuel choices.

The production model is consistent with recent evidence suggesting heterogeneity in fuel productivity and high fixed costs of fuel adoption. Lyubich, Shapiro and Walker (2018) finds that firms vary substantially in energy and CO2 productivity. These disparities in productivity are due to varying heat efficiency that different fuel-burning technologies provide (Allcott and Greenstone, 2012), energy retrofit efforts to curb energy waste (Christensen, Francisco and Myers, 2022), unobserved fuel quality (e.g., anthracite vs. bituminous coal), and intangible factors such as the ability of workers to use different fuels and management practices (Gosnell, List and Metcalfe, 2020). As in Scott (2021), I find significant fixed costs and time commitments associated with adopting natural gas. These costs encompass technological adaptations, new storage facilities, and transportation infrastructure, all microfoundations that my approach captures.

Moreover, I show that fuel-specific productivity that underlies technology differences across plants directly connects to marginal abatement cost heterogeneity. I thus contribute to a longstanding literature that estimates marginal abatement costs(Atkinson and Luo, 2023, Carlson, Burtraw, Cropper and Palmer, 2000, Culler and Mansur, 2017, Shapiro and Walker, 2018). These fuel-specific productivity differences, along with dynamic switching between fuel sets, also provide a cautionary tale against an aggregate production function to study fuel substitution, common in the integrated assessment literature (Golosov, Hassler, Krusell and Tsyvinsky, 2014, Miftakhova and Renoir, 2021), because such a production function will not be invariant to policy changes.

The production model also reflects important channels of firm responses to changes in fuel prices. Numerous empirical studies have shown that firms respond to changes in fuel cost by substituting across fuels (Ahmadi and Yamazaki, 2020, Alpino, Citino and Frigo, 2023, Andersson, 2019), but also by passing on the cost increase to consumers, thereby reducing output (Fontagné, Martin and Orefice, 2023, Ganapati, Shapiro and Walker, 2020, Gittens, 2020). While both types of responses can reduce emissions, their welfare implications are unequal, and my model provides a cohesive framework to aggregate these responses and relate them to marginal abatement costs.

To estimate this model, I contribute to the literature on production function estimation in industrial organization[3]. I make a methodological contribution by showing how to identify and estimate a dynamic production function with input-augmenting productivity and dynamic input selection. I solve this problem by drawing from the dynamic discrete choice literature with unobserved heterogeneity (Arcidiacono and Jones, 2003, Arcidiacono and Miller, 2011).

Within the production function estimation literature, my paper most closely relates to Collard-Wexler and De Loecker (2015), who studies technological change in the U.S. Steel Industry (particularly the introduction of electric arc furnaces), and Hawkins-Pierot and Wagner (2022), who estimates the energy productivity of manufacturing plants and its implication for technology lock-in. I complement the former by emphasizing the role of fuels and emissions as part of this technological change. I complement the latter by decomposing energy productivity into the relative productivity of different fuels. I show that this distinction is crucial to understanding the heterogeneous impact of a carbon tax.

The rest of the paper is structured as follows: Section 2.2 presents an overview of the Indian steel dataset. Section 2.3 presents some key evidence of plant-level decisions that motivate modeling choices. Section 2.4 presents the model in detail. Section 2.5 presents identification results for the production function. Section 2.6 presents identification results for the dynamic discrete choice model. Section 2.7 presents the main estimation results. Finally, Section 2.8 presents the results of the counterfactual experiments.

2.2 Data

I use longitudinal data on prices and quantities of all inputs and outputs from Indian steel establishments, which I link to data on India's natural gas pipeline network. The panel comes from the Indian Survey of Industries (ASI) and covers 2009-2016. It is a restricted-use dataset that covers all manufacturing establishments with at least 100 workers and a representative sample of establishments with fewer than 100 workers. The sample is stratified at various levels, including number of workers and location. More details on sampling rules, including changes over time, can be found in Appendix A.1. The ASI contains information on prices and quantities of Coal, Oil, Electricity, and Natural Gas, which I convert to million British thermal units (mmBtu) using standard scientific calculations from the U.S. Environmental Protection Agency

[3]See footnote[2] for details on this literature.

(EPA, 2022). I follow standard practices by removing the 1% left and right tails of plant-level inputs and output by industry.[4]

To convert nominal into real units, I follow Harrison, Hyman, Martin and Nataraj (2016) by deflating output with industry-specific wholesale price indices (WPI), labor with the consumer price index (CPI), intermediate materials with the aggregate wholesale price index, labor with the consumer price index (CPI), and capital stock with an India-specific capital deflator from the Penn World Table (Feenstra, Inklaar and Timmer, 2015).

Emissions To get establishment-level measures of greenhouse gas emissions, I convert units of potential energy (mmBtu) of each fuel into metric tons of carbon dioxide equivalent (CO_{2e}). During combustion, each mmBtu of fuel releases some quantity of carbon dioxide CO_2, methane CH_4, and nitrous oxide N_2O into the atmosphere, which varies by industry based on standard practices in India (Gupta, Biswas, Janakiraman and Ganesan, 2019, Annexure 3). I then convert emissions of these three chemicals into carbon dioxide equivalent (CO_{2e}) using the Global Warming Potential method (GWP, see Appendix A.2).

2.3 Facts About Emissions and Fuels in India

Using this data, I highlight facts about fuel usage and carbon emissions that motivate my choice of India's Steel sector to conduct this analysis and influence modeling decisions.

Fact 1: High Pollution Levels from Indian Steel Establishments

In Table 2.1, I show that total annual greenhouse gas emissions from Indian Steel plants average 25 million tons of CO_{2e}, accounting for 31% of annual emissions in Indian manufacturing (Dhar, Pathak and Shukla, 2020). This high emission level is attributed to the sizeable aggregate share of coal as part of the energy mix, averaging 70%. This share is significantly larger than in other Indian manufacturing industries and larger than in Steel manufacturing abroad. Indeed, switching from coal to gas has contributed significantly to the manufacturing clean-up in developed economies (Rehfeldt, Fleiter, Herbst and Eidelloth, 2020).

Fact 2: Indian Steel Establishments Use Different Fuel Sets

Steel-producing plants use different fuel sets, and most fuel sets include oil and electricity. Most of the variation in fuel sets thus comes from whether plants use coal, natural gas, neither,

[4] Such outliers are typically due to reporting errors and are inconsistent with a wide range of official statistics (Bollard, Klenow and Sharma, 2013).

Industry	Annual Average Number of Plants	Average Annual Revenue by Plant (Million USD)	Annual Average Emissions (Thousand tons CO_{2e})	Aggregate Energy Input Share	Aggregate Coal Fuel Share
Steel	1,077	19.41	29.34	0.13	0.72
Other	33,726	6.18	8.52	0.13	0.37

Table 2.1: Descriptive Statistics for Steel Manufacturing (2009-2016)

Note: The energy input share is calculated as the aggregate spending on energy by industry as a fraction of total spending on labor, materials, and energy. It is then averaged across years. Similarly, the coal fuel share is calculated as the aggregate share of coal (in mmBtu) relative to other fuels in each industry, averaged across years.

or both. see Table 2.2. There are multiple reasons for this heterogeneity. For example, plants can use different furnaces to turn iron ore into steel. Blast furnaces combined with basic oxygen furnaces rely on coke (coal) as a primary fuel; electric arc furnaces can use any fuel combination to generate high-power electricity, which is discharged through an electric arc to melt either steel scrap or sponge iron. Sponge iron is created from iron ore in direct reduced iron furnaces powered by natural gas.

	Percentage of Plants	Output Share
Oil, Electricity	51.3	44.9
Oil, Electricity, Coal	19.3	21.2
Oil, Electricity, Gas	10.8	25.7
Oil, Electricity, Coal, Gas	7.4	3.4
Other	11.1	4.8

Table 2.2: Distribution of Fuel Sets Across Steel Plants

Notes: This table shows the distribution of fuel sets across plants. The “Other” category comprises any other combinations of the same four fuels. Variation in fuel sets is not driven by variation in the variety of steel produced. In Appendix A.3, I show that a similar distribution exists within different varieties of steel produced.

Many reasons explain why plants use different fuel sets, some of which have geographical underpinnings. In Figure 2.1, I document a concentration of coal usage in Eastern and Southern India. Many states in Eastern India form a region colloquially known as the “Steel Belt” due to the prevalence of coal and iron ore mines. Plants with coal-powered Blast furnaces tend to be located near these mines. Moreover, this region lacks critical infrastructure, such as natural gas pipelines for plants to adopt technologies that rely on cleaner fuels.

This heterogeneity in fuel sets has two other noteworthy implications. First, burning coal releases more CO_{2e} than burning natural gas. Second, plants with more fuels have access to additional margins of substitution, which is relevant because many fossil fuels are susceptible to price volatility owing to global supply and demand fluctuations. For instance, oil and natural gas prices fell due to booming U.S. shale oil production and excess supply from emerging market economies (EMEs) such as Saudi Arabia in 2014. See Figure 2.5.

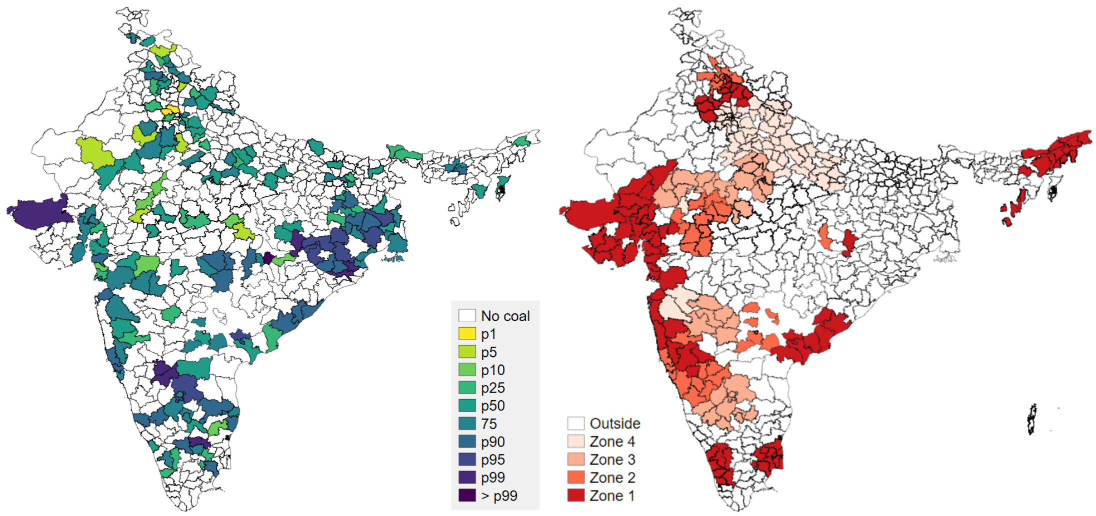

(a) Coal Consumption by District (b) Natural Gas Pipeline Network (2016)

Figure 2.1: Spatial Variation in Fuel Usage and Key Infrastructures

Notes: The map on the left shows the distribution of coal usage across Indian districts. Darker areas correspond to districts with more coal consumption. The map on the right shows the distribution of access to natural gas pipelines across districts. Natural gas pipeline transportation tariffs are organized by zones, where each zone corresponds to 250km segments along a pipeline. Zone 1 is the closest to the source of the pipeline. Moreover, at the beginning of the sample period, the natural gas pipeline network was not developed in the South, and it was less developed in the Delhi area.

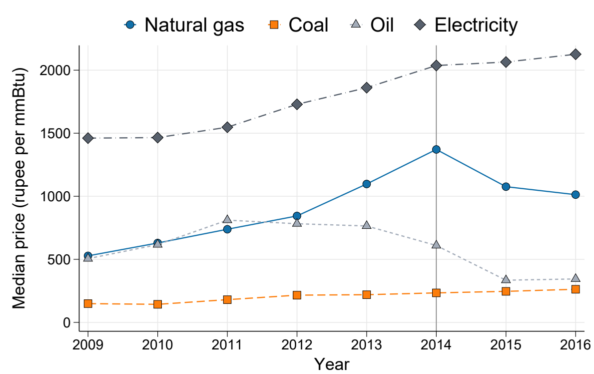

Figure 2.2: Median Fuel Prices (rupees/mmBtu)

Fuel substitution can serve as a means of adjustment for plants to insure themselves against fuel price variation. Both the quantity and fuel share of coal increased prior to 2014. However, coal usage started to decline after 2014 following the crash in the price of oil and natural gas. See Figure 2.3 and 2.4. Electricity prices have been steadily increasing in India, and the share of electricity in steel plants' energy mix has been steadily decreasing.

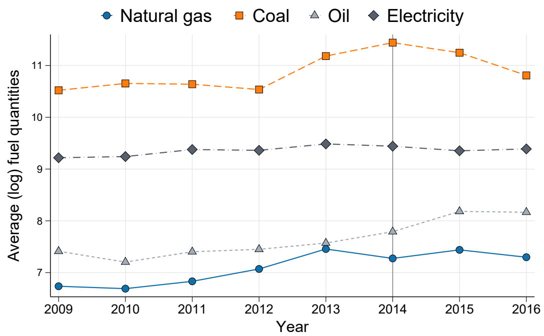

Figure 2.3: Average (log) Fuel Quantities

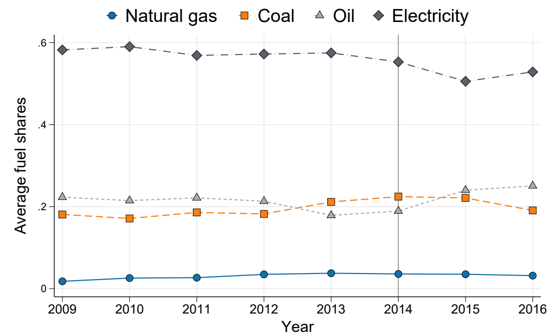

Figure 2.4: Average Fuel Shares

Notes: These figures show the average quantities of each fuel (in logs) and the average within-plant fuel share, respectively. Fuel quantities are measured in mmBtu for each fuel, and these figures are constructed from steel-producing plants only.

Fact 3: Indian Steel Establishments Often Switch Between Fuel Sets

Not only is there heterogeneity in fuel sets but also a significant prevalence of switching between fuel sets, which I define as occurring when a plant uses a different combination of fuels between two years (e.g., from oil and electricity to oil, electricity, and gas). On average, 15% of plants add a new fuel, and 15% drop an existing fuel from their set every year. Moreover, 40% of unique plants add and drop fuel at least once in the sample. There are many reasons why plants switch between fuel sets. The development of new technologies may increase the productivity of some fuels. Electric arc furnaces are more efficient at using the heating potential of underlying fuels than blast furnaces (Worrell, Bernstein, Roy, Price and Harnisch, 2009). Large and persistent fuel price shocks incentivize plants to readjust their input mix. Expanding transportation infrastructures, particularly pipeline networks, decreases fixed costs and eases access to new fuels (Scott, 2021).

	Adds New Fuel	Drop Existing Fuel	Both Add and Drop
Yes (%)	39.4	39.6	26.0

Table 2.3: Fraction of Unique Plants That Add and Drop a Fuel at Least Once

Notes: This table shows the fraction of unique plants that add a fuel at least once in the sample, and similarly for plants that drop a fuel at least once. This underestimates the prevalence of fuel switching because plants are only observed between 2009 and 2016.

Fact 4: Larger Fuel Set Is Associated With Higher Productivity

Adopting a new fuel is associated with substantial fixed costs: acquiring and installing new equipment such as furnaces, transportation infrastructure, and storage facilities. With fixed costs, it is natural to expect larger and more productive plants to use more fuels because they

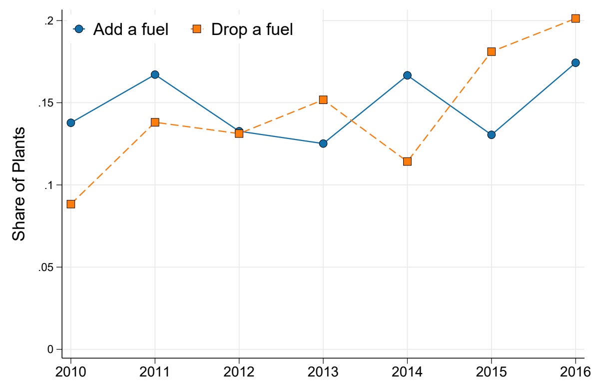

Figure 2.5: Fuel Set Switching Across Years

Notes: This figure shows the fraction of plants that add a new fuel or drop an existing fuel between year t and $t-1$, starting from 2010 all the way until 2016.

have marginally more to gain from fixed investments. This is what I find. As plants produce more output per worker, which can be interpreted as a proxy for productivity, they tend to have a larger fuel set. See Figure 2.6

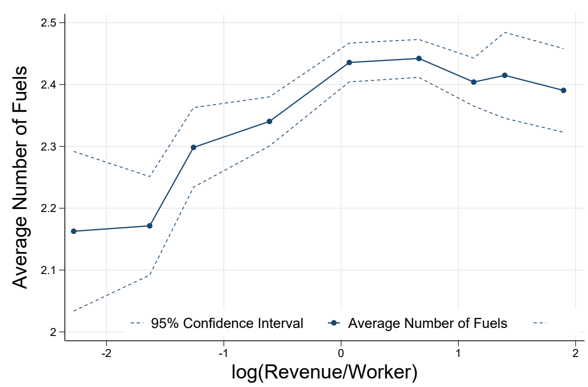

Figure 2.6: Number of Fuels, by Revenue per Worker

Notes: This figure reports the average number of fuels across percentiles of the residualized revenue per worker distribution. Revenue per worker is residualized by taking out year-fixed effects.

In Table 2.4, I show that a 1% increase in Revenue per Worker decreases the probability of having one or two fuels and increases the probability of having 3 or 4 fuels. A positive gradient between the number of fuels and productivity coupled with high fixed costs of new fuel adoption may lock unproductive plants into their current fuel set. A similar technological lock-in has been previously documented in manufacturing by Hawkins-Pierot and Wagner (2022) and decreases policy effectiveness at incentivizing switching to cleaner fuels.

Increasing Revenue/Worker by 1%		
1 Fuel	-0.009***	(0.001)
2 Fuels	-0.025***	(0.004)
3 Fuels	0.023***	(0.004)
4 Fuels	0.010***	(0.002)
Year Fixed Effects	Yes	
N	8,583	

Table 2.4: Relationship Between Revenue/Worker and Number of Fuels

Notes: This table reports marginal effects of an ordered logit regression of the number of fuels against log(Revenue/Productivity). The point estimates are the percentage increase (decrease) in the probability of having each number of fuels for a 1% increase in Revenue per Worker.

2.4 Model

Consistent with these facts, I develop and estimate a rich dynamic production model to quantify establishments' fuel choices. Each period, plants have access to a set of fuels from a combination of oil, natural gas, coal, and electricity. Fuels are combined to produce energy that goes into the outer nest of production. Plants can choose to change fuel sets across periods in a dynamic discrete choice framework. There are fixed costs for adding new fuels and salvage values from dropping existing ones. I first present the production structure for a given plant in a static setting and then consider inter-temporal decisions. Throughout the exposition, subscript i refers to a plant, and t refers to a year.

2.4.1 Production Model

There are two levels of production, which correspond to two nests. The outer nest is a standard CES production function and features Hicks-neutral productivity z_{it}, labor L_{it}, capital K_{it}, intermediate inputs M_{it}, and energy E_{it}. Following Grieco et al. (2016), the production function is explicitly normalized around the geometric mean of each variable $\overline{X} = \left(\prod_{i=1}^{n}\prod_{t=1}^{T} X_{it}\right)^{\frac{1}{nT}}$.[5]

$$\frac{Y_{it}}{\overline{Y}} = z_{it}\left(\alpha_k\left(\frac{K_{it}}{\overline{K}}\right)^{\frac{\sigma-1}{\sigma}} + \alpha_L\left(\frac{L_{it}}{\overline{L}}\right)^{\frac{\sigma-1}{\sigma}} + \alpha_M\left(\frac{M_{it}}{\overline{M}}\right)^{\frac{\sigma-1}{\sigma}} + \alpha_E\left(\frac{E_{it}}{\overline{E}}\right)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\eta\sigma}{\sigma-1}} \tag{2.1}$$
$$s.t. \quad \alpha_L + \alpha_K + \alpha_M + \alpha_E = 1$$

[5]All CES functions are either implicitly or explicitly normalized around a point (León-Ledesma, McAdam and Willman, 2010). I chose the geometric mean as a normalization point to be consistent with the literature.

Where $\sigma \geq 0$ is the elasticity of substitution between inputs, and $\eta > 0$ is the returns to scale. In the outer nest, plants choose input quantities given input prices, which include energy, E_{it}. Given the current fuel set $\mathcal{F}_{it} \subseteq \mathbb{F} = \{\text{oil, gas, coal, elec}\}$, plants combine all fuels available to produce a quantity of energy E_{it} in the inner nest of production:

$$\frac{E_{it}}{\overline{E}} = \left(\sum_{f \in \mathcal{F}_{it}} \left(\psi_{fit} \frac{e_{fit}}{\overline{e_f}} \right)^{\frac{\lambda-1}{\lambda}} \right)^{\frac{\lambda}{\lambda-1}} \tag{2.2}$$

e_{fit} refers to the quantity of fuel f for plant i in year t. p_{fit} and ψ_{fit} are the corresponding fuel price and productivity, respectively. The fuel-specific productivity terms are novel; they allow for flexible variation in input usage and heterogeneity in fuel substitution. Plants specialize in fuels that they can use more efficiently, which means that different plants are affected differently by changes in fuel prices. This heterogeneity is especially relevant in the context of a carbon tax, which raises the price of dirty fuels more than clean fuels. For example, plants specializing in dirty fuels such as coal will bear a disproportionate share of the tax burden. While initially unobserved, I recover fuel productivity for each plant each year by exploiting profit maximization.

Moreover, allowing plants to have different fuel sets $\mathcal{F}_{it}$ and allowing plants to switch between them is novel; plants with a larger fuel set have more substitution possibilities when facing changes in fuel prices. This option value creates another layer of heterogeneity in response to carbon taxation. Larger and more productive plants with a larger fuel set will have an easier time substituting out of a carbon tax than smaller plants with smaller fuel sets.

These two novel features are a significant departure from the literature, where most previous papers that estimate a production function with fuels do not allow for fuel-specific productivity and do not allow for fuel sets to vary within the same production function (Atkinson and Halvorsen, 1976, Hyland and Haller, 2018, Joskow and Mishkin, 1977, Ma et al., 2008, Pindyck, 1979). More recently, Hawkins-Pierot and Wagner (2022) allowed for the productivity of the total energy bundle to vary across plants. While this allows for heterogeneity in the substitution between energy and other inputs, it does not capture salient features of fuel consumption and differential responses to fuel price changes.

The elasticity of substitution between fuels λ plays a crucial role in this model. It determines the option value a plant would get by expanding its fuel set $\mathcal{F}_{it}$. As long as fuels are gross substitutes ($\lambda > 1$), there is an option value (gains from variety) from having more fuels due to the additional substitution margin each fuel provides. However, the lower λ is, the larger the option value. A lower λ implies that marginal products from a given fuel decrease faster with quantity, so there are larger marginal gains from adding a new fuel[6]. Next, I show how plants compete and set prices.

2.4.2 Static Decisions

Plants produce different output varieties and engage in monopolistic competition.

I make this assumption for two reasons. First, the Indian steel industry features the most plants out of all Indian heavy manufacturing industries, and the industry faces fierce competition from China, the world's largest steel producer. It is thus likely to be more competitive than typical oligopolistic industries like Cement. Second, steel plants produce a wide variety of steel products. There are 404 varieties produced by plants in the ASI. *Ferrous products from direct reduction of iron ore* are the most common variety with a 5.5% market share.

On the demand side, there is a representative consumer with quasi-linear utility over the total output produced in a given period Y_t and an outside good Y_{0t}. Steel consumption is widespread in India, and most demand comes from housing construction, infrastructure, and automobiles. Total output is produced by aggregating all the varieties with standard Dixit-Stiglitz preferences across varieties. Given a mass of N_t operating plants, income I_t, and an aggregate demand shock e^{Γ_t}, the representative consumer solves:

$$\begin{aligned} \max_{\{Y_{it}\}_{i=1}^{N_t}, Y_{0t}} \ \mathbb{U} &= Y_{0t} + \frac{e^{\Gamma_t}}{\theta}\left(\frac{1}{N_t}\int_{\Omega_i}(N_t Y_{it})^{\frac{\rho-1}{\rho}}di\right)^{\frac{\theta\rho}{\rho-1}} \\ s.t. \quad Y_{0t} &+ \int_{\Omega_i} P_{it}Y_{it}di \leq I_t \end{aligned} \tag{2.3}$$

Where $\rho > 1$ is the elasticity of substitution between varieties, and $\theta \in (0, 1)$ indexes the substitution between consumption of the differentiated varieties and the outside good. Following Helpman and Itskhoki (2010), I restrict $\theta < \frac{\rho-1}{\rho}$, which ensures that output varieties are more substitutable between each other than with the outside good. These quasi-linear CES preferences were first proposed by Helpman and Itskhoki (2010) and provide analytical convenience for welfare evaluation. Quasi-linear preferences are standard in the literature on externality taxation (Fowlie et al., 2016) and allow researchers to use the social cost of carbon (SCC), which

[6]This option value is similar to the concept of gains from variety in the trade literature investigating the composition of intermediate inputs (Broda and Weinstein, 2006, Ethier, 1982, Goldberg, Khandelwal, Pavcnik and Topalova, 2010, Kasahara and Rodrigue, 2008, Ramanarayanan, 2020, Romer, 1990)

expresses the net present value of expected future damages from carbon emissions in dollars.[7] Externality damages thus affect consumption of the outside good by varying aggregate income and thus directly affect consumer surplus. Solving the representative consumer's problem in (2.3) yields the following downward-sloping demand for each variety Y_{it}, which I augment with an ex-post idiosyncratic demand shock $e^{\epsilon_{it}}$:

$$Y_{it} = \frac{e^{\tilde{\Gamma}_t}}{N_t} P_{it}^{-\rho} P_t^{\frac{\rho(1-\theta)-1}{1-\theta}} e^{\epsilon_{it}} \tag{2.4}$$

Where $e^{\tilde{\Gamma}_t} = e^{\Gamma_t \frac{1}{1-\theta}}$ and $P_t = \left(\frac{1}{N_t}\int P_{it}^{1-\rho} di\right)^{\frac{1}{1-\rho}}$ is the CES aggregate price index across all varieties. Detailed derivations can be found in Appendix B.1.

Plants choose inputs to maximize static profits

Given a set of fuels $\mathcal{F}_{it} \subseteq \mathbb{F}$, technological constraints, inverse demand, and all input prices, the plant's problem is static.[8] To avoid notation clutter, I will define $\tilde{X}_{it} \equiv \frac{X_{it}}{\overline{X}}$ for normalized quantities and $\tilde{p}_{xit} \equiv p_{xit}\overline{X}$ for normalized prices from now on.

$$\begin{aligned}
&\max_{K_{it}, M_{it}, L_{it}, \{e_{fit}\}_{f\in\mathcal{F}_{it}}} \left\{P_{it}(Y_{it})Y_{it} - w_t L_{it} - r_{kt}K_{it} - p_{mit}M_{it} - \sum_{f\in\mathcal{F}_{it}} p_{fit}e_{fit}\right\} \\
s.t. \quad &\tilde{Y}_{it} = z_{it}\left[\alpha_K \tilde{K}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_L \tilde{L}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_M \tilde{M}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_E\Big(\sum_{f\in\mathcal{F}_{it}} (\psi_{fit}\tilde{e}_{fit})^{\frac{\lambda-1}{\lambda}}\Big)^{\frac{\lambda}{\lambda-1}\frac{\sigma-1}{\sigma}}\right]^{\frac{\eta\sigma}{\sigma-1}} \\
&P_{it}(Y_{it}) = \left(\frac{e^{\tilde{\Gamma}_t}}{N_t Y_{it}}\right)^{\frac{1}{\rho}} P_t^{\frac{1+\rho(\theta-1)}{(\theta-1)\rho}}
\end{aligned}$$

The nested structure of production is such that it can be expressed in two stages:

1. Plants choose fuel quantities to minimize the cost of producing energy (inner nest):

Given a fuel set $\mathcal{F}_{it}$ and fuel prices, plants find the combination of fuels that minimizes the cost

[7] This is the approach typically taken in applied microeconomics. However, an alternative approach in macroeconomics relies on integrated assessment models (IAM) to explicitly study the dynamic relationship between aggregate emissions and the concentration of CO_2 in the atmosphere, which affects future aggregate output in various ways. See Golosov et al. (2014), Hassler et al. (2019, 2020), Nordhaus (2008).

[8] I derive the decision of plants under the assumption that plants flexibly rent capital with a unit cost of capital r_{kt}. While I use this assumption to reduce the computational burden in the dynamic discrete choice model of fuel sets, I do not need nor use this assumption to estimate the production function.

of producing a given unit of energy. Fuel prices in mmBtu are observed and vary across plants and year:

$$\min_{\{e_{fit}\}_{f\in\mathcal{F}_{it}}} \Big\{ \sum_{f\in\mathcal{F}_{it}} p_{fit} e_{fit} \Big\} \quad s.t. \quad \tilde{E}_{it} = \Big(\sum_{f\in\mathcal{F}_{it}} (\psi_{fit}\tilde{e}_{fit})^{\frac{\lambda-1}{\lambda}} \Big)^{\frac{\lambda}{\lambda-1}} \tag{2.5}$$

The achieved minimum of this problem is an energy cost function $C(\tilde{E}_{it})$ that satisfies:

$$\begin{aligned} C(\tilde{E}_{it}) &= \Big(\sum_{f\in\mathcal{F}_{it}} (\frac{\tilde{p}_{fit}}{\psi_{fit}})^{1-\lambda} \Big)^{\frac{1}{1-\lambda}} \tilde{E}_{it} \\ &= p_{\tilde{E}_{it}} \tilde{E}_{it} = \sum_{f\in\mathcal{F}_{it}} p_{fit} e_{fit} \end{aligned}$$

Where the unobserved price of realized energy $\tilde{p}_{E_{it}}$ corresponds to a CES price index in fuel prices over productivity. Constant returns in the energy production function imply that the marginal cost of realized energy is the price of realized energy and is constant $MC(\tilde{E}_{it}) = p_{\tilde{E}_{it}}$.

2. Plants choose inputs other than fuels to maximize profit (outer nest):

Given a cost-minimizing allocation of fuels that produce a quantity of energy, plants pay a price p_{Eit} for each unit of energy. They take this price as given when choosing the quantity of energy because p_{Eit} is only a function of the optimal *relative* allocation of fuels, not the scale of energy. Then, at the beginning of each period, plants start with a set of fuels $\mathcal{F}_{it} \subseteq \mathbb{F}$, observe their Hicks-neutral productivity z_{it}, productivity for each fuel $\{\psi_{fit}\}_{f\in\mathcal{F}_{it}}$, and all input prices $\{w_{it}, r_{kit}, p_{mit}, \{p_{fit}\}_{f\in\mathcal{F}_{it}}\}$. Together with the years of production, these form a set of state variables s_{it}. Given these state variables, plants maximize profits, which yield a period profit function $\pi(\mathbf{s}_{it}, \mathcal{F}_{it})$.

$$\begin{aligned} \pi(\mathbf{s}_{it}, \mathcal{F}_{it}) = &\max_{K_{it},M_{it},L_{it},E_{it}} \Big\{ \Big(\frac{e^{\Gamma_t}}{N_t}\Big)^{\frac{1}{\rho}} P_t^{\frac{1+\rho(\theta-1)}{(\theta-1)\rho}} Y_{it}^{\frac{\rho-1}{\rho}} - w_t L_{it} - r_{kt} K_{it} - p_{mit} M_{it} - p_{Eit} E_{it} \Big\} \\ s.t. \quad \tilde{Y}_{it} &= z_{it} \Big[\alpha_K \tilde{K}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_L \tilde{L}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_M \tilde{M}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_E \tilde{E}^{\frac{\sigma-1}{\sigma}} \Big]^{\frac{\eta\sigma}{\sigma-1}} \end{aligned} \tag{2.6}$$

2.4.3 Inter-Temporal Fuel Set Choices

Every period, plants take expectations over the evolution of state variables and choose a fuel set for the next period $\mathcal{F}'$ to maximize expected discounted lifetime profits:

$$V(\mathbf{s}_{it}, \mathcal{F}_{it} \in \mathbb{F}) = \max_{\mathcal{F}'} \Big\{ \underbrace{\pi(\mathbf{s}_{it}, \mathcal{F}_{it})}_{\text{static profits}} - \underbrace{\mathcal{K}(\mathcal{F}' \mid \mathcal{F}_{it}, s_{it}) + \sigma_{\epsilon}\epsilon_{\mathcal{F}'it}}_{\text{fixed switching costs}} + \underbrace{\beta\, \mathbb{E}[V(\mathbf{s}_{it+1}, \mathcal{F}') \mid s_{it}]}_{\text{continuation value}} \Big\}$$

Where $\mathcal{K}(\mathcal{F}' \mid \mathcal{F}_{it}, s_{it})$ is the net cost of switching from fuel set $\mathcal{F}$ to $\mathcal{F}'$, and $\epsilon_{\mathcal{F}'_{it}}$ captures idiosyncratic shocks to these switching costs. I allow fuel set switching costs to vary by plant size (proxied by Hicks-neutral productivity z_{it}) and whether a plant is in a district d that has access to natural gas pipelines[9]:

$$\mathcal{K}(\mathcal{F}' \mid \mathcal{F}_{it}, s_{it}) = k(\mathcal{F}' \mid \mathcal{F}_{it}, d_{it}) + \gamma \ln z_{it}$$

The switching cost function $k(\mathcal{F}' \mid \mathcal{F}_{it}, d_{it})$ is composed of two types of arguments. First, there are fixed costs of adding a fuel κ_f. Second, there are salvage values of dropping a fuel γ_f that plants obtain by selling old capital. Since 90% of plants in the dataset always use electricity and oil, I assume that the choice set of plants is as follows, where e = electricity, o = oil, g = gas, c = coal: $\mathbb{F} = \Big\{(oe); (oge); (oce); (ogce)\Big\}$ and restrict the sample accordingly. In the next section, I show how this model can be estimated.

2.5 Identification of the Production Function

The model is estimated in four steps using a novel combination of methods. First, I estimate the demand elasticity by using the shift-share instrument of Ganapati et al. (2020) as a cost-shifter. Second, I adapt the method of Grieco et al. (2016) in order to estimate the outer production function in the presence of an unobserved input (energy). Third, I estimate the energy production function following recent developments in production function estimation with input-augmenting productivity (Demirer, 2020, Zhang, 2019), combined with dynamic panel

[9] Plant size is endogenous, but a Ceteris paribus increase in z_{it} increases the scale of a plant's operation. Scott (2021) shows that proximity to the natural gas pipeline network decreases the fixed cost of adding natural gas. Plants too far from the pipeline network can use liquified natural gas (LNG) but need access to a gasification terminal, which can be very costly.

techniques (Blundell and Bond, 1998, 2000, 2023). Fourth, I estimate fixed costs in a dynamic discrete choice framework in the presence of unobserved heterogeneity following Arcidiacono and Jones (2003), which allows me to capture systematic differences in fuel productivity for fuels that plants are not currently using.

2.5.1 Identification of Outer Production Function

In the outer nest, the main unobserved quantity that departs from standard models is realized energy $\tilde{E}_{it}$. In contrast to the heating potential of fuels, energy is the output of combining different fuels, which is unobserved. I adapt the estimation method proposed by Grieco et al. (2016) to uniquely recover the price and quantity of energy when other flexible inputs are observed under the assumption that plants are price-takers in the input market.[10] The key to this method relies on using relative first-order conditions to map observed expenditure shares to unobserved input quantity shares. To see this, one can look at the ratio of first-order conditions for labor and energy from profit maximization in equation 2.6 and rearranging:

$$\underbrace{\frac{w_{it}L_{it}}{p_{E_{it}}E_{it}}}_{\text{Expenditure ratio}} = \frac{\alpha_L}{\alpha_E}\Big(\underbrace{\frac{L_{it}/\overline{L}}{E_{it}/\overline{E}}}_{\text{Quantity ratio}}\Big)^{(\sigma-1)/\sigma} \tag{2.7}$$

Given production function parameters, $\frac{E_{it}}{\overline{E}}$ can be recovered from (2.7) because I observe ex-

[10]The assumption of price-taking in the input market allows for unobserved variation in input prices (the main motivation underlying the Grieco et al. (2016) paper), which could be related to plant size, productivity, location, and any other state variables. However, this assumption rules out quantity discounts.

$\mathbb{F}$	**oe**	**oge**	**oce**	**ogce**
oe	0	κ_g	κ_c	$\kappa_g + \kappa_c$
oge	$-\gamma_g$	0	$-\gamma_g + \kappa_c$	κ_c
oce	$-\gamma_c$	$-\gamma_c + \kappa_g$	0	κ_g
ogce	$-\gamma_g - \gamma_c$	$-\gamma_c$	$-\gamma_g$	0

Table 2.5: Fixed Switching Cost Matrix — $k(\mathcal{F}' \mid \mathcal{F})$

Notes: Rows correspond to fuel sets today $\mathcal{F}$, whereas columns correspond to fuel sets next period $\mathcal{F}'$. I assume that fixed costs and salvage values for coal are the same across districts. However, fixed costs and salvage values for natural gas vary by plants' proximity to the natural gas pipeline network in a binary fashion. I define $d = 0$ if plants have no access to pipelines and $d = 1$ if plants have access to pipelines. Then $\kappa_g = \kappa_{g1}$ if $d = 1$ and $\kappa_g = \kappa_{g0}$ if $d = 0$, and likewise for γ_g. I define plants as having access to pipelines if they are located in a district in which a pipeline directly passes or in a district immediately adjacent to a district in which a pipeline passes.

penditures for both inputs (recalling that energy expenditure is the sum of fuel expenditures from the energy production function: $p_{E_{it}}E_{it} = \sum_{f\in\mathcal{F}_{it}} p_{fit}e_{fit}$) and I observe the quantity of labor. Identification of $\tilde{E}_{it}$ comes from variation in the relative price of labor to energy, which induces variation in the expenditure ratio that isn't one-for-one with relative prices. For a given σ, observed variation in spending on energy $S_{E_{it}}$, spending on labor $S_{L_{it}}$ and the quantity in labor L_{it} implies a unique quantity of realized energy by the optimality condition between both inputs. Only when $\sigma = 1$ (Cobb-Douglas), the percentage change in relative prices is always offset by an equivalent percentage change in expenditure shares, such that expenditure shares are constant.

$$\frac{E_{it}}{\overline{E}} = \Big(\frac{p_{eit}E_{it}}{w_{it}L_{it}}\Big)^{\frac{\sigma}{\sigma-1}}\Big(\frac{\alpha_L}{\alpha_E}\Big)^{\frac{\sigma}{\sigma-1}}\frac{L_{it}}{\overline{L}} \tag{2.8}$$

Thus, in this setting, one can identify production parameters by replacing $\tilde{E}_{it}$ for (2.8) in the production function and exploiting first-order conditions to control for the transmission bias from unobserved hicks-neutral productivity z_{it} to observed inputs. This method is also used by Doraszelski and Jaumandreu (2013, 2018). I also use the same method to control for unobserved price dispersion in the bundle of material inputs:

$$\frac{M_{it}}{\overline{M}} = \Big(\frac{p_{mit}M_{it}}{w_{it}L_{it}}\Big)^{\frac{\sigma}{\sigma-1}}\Big(\frac{\alpha_L}{\alpha_M}\Big)^{\frac{\sigma}{\sigma-1}}\frac{L_{it}}{\overline{L}}$$

The main dependent variable is revenues, where $e^{u_{it}}$ is an unobserved iid shock that is meant to capture measurement error and unanticipated demand & productivity shocks to the plant (Klette and Griliches, 1996). Detailed derivations of the estimating equation can be found in Appendix C.1. Taking logs of revenues yields the main estimating equation:

$$\ln R_{it} = \ln\frac{\rho}{\rho-1} + \ln\frac{1}{\eta} + \ln\Bigg[w_{it}L_{it}\Big(1+\frac{\alpha_K}{\alpha_L}\Big(\frac{K_{it}/\overline{K}}{L_{it}/\overline{L}}\Big)^{\frac{\sigma-1}{\sigma}}\Big) + p_{mit}M_{it} + p_{eit}E_{it}\Bigg] + u_{it} \tag{2.9}$$

The main parameters of interest are the elasticity of substitution (σ), the elasticity of demand (ρ), and the returns to scale (η) in (2.9). While the elasticity of substitution is identified from observed variation in the capital-to-labor ratio, the elasticity of demand/markup is not separately identified from the returns to scale. This is a standard problem with revenue production functions, whereby the curvature in the revenue function is driven by both technology (returns to scale) and market power (markup). Fortunately, I observe output prices and quantities, and I use exogenous cost shifters to recover the elasticity of demand ρ in Section 2.5.1. Lastly, since $\tilde{E}_{it}$ and $\tilde{M}_{it}$ were factored out of the production function, the main estimating equation (2.9) does not recover α_E and α_M. To recover α_E and α_M, I take the geometric mean of relative first-order conditions in equation (2.7) for labor relative to energy and for labor relative to intermediate materials.[11]

$$\begin{aligned} \overline{wL}/\overline{p_E E} = \frac{\alpha_L}{\alpha_E}; \qquad \overline{wL}/\overline{p_m M} = \frac{\alpha_M}{\alpha_E} \\ \alpha_K + \alpha_L + \alpha_M + \alpha_E = 1 \end{aligned} \tag{2.10}$$

Then, I estimate (2.9) subject to (2.10) with non-linear least squares.[12]

Estimating Elasticity of Demand

To separate the demand elasticity ρ from the returns to scale η in the estimating equation (2.9), I estimate demand from observed output prices and quantities using the demand equation (2.4).

$$\ln Y_{it} = \Lambda_t - \rho \ln P_{it} + \epsilon_{it}, \tag{2.11}$$

where $\Lambda_t = \tilde{\Gamma}_t + \ln\left(\frac{1}{N_t}\right) + \frac{\rho(1-\theta)-1}{1-\theta} \ln P_t$ contains both the unobserved aggregate output price index P_t and aggregate demand shocks $\tilde{\Gamma}_t$. Due to standard simultaneity bias, the elasticity of demand ρ is not identified from price and quantity data alone. To solve this issue, I instrument output prices with a Barktik style shift-share cost shifter proposed by Ganapati et al. (2020) and used by Hawkins-Pierot and Wagner (2022). The instruments have two components: an exogenous shock to aggregate fuel prices (the shift) and a pre-shock variation in exposure to aggregate fuel prices by Indian States (the share):

[11] This is the convenience given by the geometric mean normalization of the CES. However, any other normalization would work but would require more algebra to recover the distribution parameters.

[12] Consistency of the parameters is shown by Grieco et al. (2016) using the first-order conditions of the NLLS objective function as moment conditions.

$$z_{s,t} = \left[\overline{p}_{-s,t,f} * \sigma_{s,2008,f}\right], \quad f \in \{\text{coal, gas, oil}\}$$

$\overline{p}_{-s,t,f}$ is the average price (leaving out state s) of fuel f in year t and acts as an exogenous shock to production cost. This is because much of aggregate fuel price variation stems from worldwide demand and supply variation induced by geopolitical turmoil, aggregate technological evolution, and growth. $\sigma_{s,2008,f}$ is the pre-sample aggregate share of fuel f used to generate electricity in state s. Any variation in the price of a fuel will induce more variation in electricity prices in states that use more of that fuel to generate electricity. This creates exogenous variation in exposure to aggregate fuel price shocks since all plants use electricity as an input. Moreover, the shares are taken in 2008 (before the sample starts) and are thus unaffected by shocks to fuel prices.

For the remaining parts of the demand equation, the aggregate output price index P_t is part of the year-fixed effect in equation (2.11). It is endogenously determined by the elasticity of demand ρ. I first estimate demand using year dummies Λ_t and then solve for the price index ex-post given the estimate of $\hat{\rho}$, observed output prices P_{it} and the number of plants N_t, $P_t = \left(\frac{1}{N_t}\sum_{i=1:N_t} P_{it}^{1-\hat{\rho}}\right)^{1/(1-\hat{\rho})}$. I then separately recover the elasticity of the outside good θ from the aggregate demand shifter $\tilde{\Gamma}_t$ in a simple time series regression of the year dummies Λ_t on the output price index and a constant.

2.5.2 Identification of Inner Production Function for Energy

The energy production function in equation (2.2) can be rewritten by factoring out the productivity of a fuel that plants always use, such as electricity, and redefining the productivity of all other fuels relative to electricity, $\tilde{\psi}_{fit} = \frac{\psi_{fit}}{\psi_{eit}}$:

$$\tilde{E}_{it} = \psi_{eit}\left(\sum_{f\in\mathcal{F}_{it}}\left(\tilde{\psi}_{fit}\frac{e_{fit}}{\overline{e}_f}\right)^{\frac{\lambda-1}{\lambda}}\right)^{\frac{\lambda}{\lambda-1}} \tag{2.12}$$

At this point, I have an estimate of the quantity and price of energy, $(\hat{E}_{it}, p_{\hat{E}_{it}})$ from the previous step, fuel quantities, $\{e_{fit}\}_{f\in\mathcal{F}_{it}}$, and fuel prices: $p_{\hat{E}_{it}} = \frac{S_{Eit}}{\hat{E}_{it}}, \{p_{fit} = \frac{s_{fit}}{e_{fit}}\}_{f\in\mathcal{F}_{it}}$. I show how to recover the elasticity of substitution λ and all productivity terms ψ_{fit}. To do so, I rely on optimality conditions from the energy cost-minimization problem coupled with a Markovian assumption on electricity productivity. This effectively combines the dynamic panel approach of Blundell and Bond (1998, 2000) with the method proposed by Zhang (2019). As a reminder, the energy cost-minimization problem of the plant is as follows:

$$\min_{\{e_{fit}\}_{f\in\mathcal{F}_{it}}} \sum_{f\in\mathcal{F}_{it}} p_{fit}e_{fit} \qquad s.t. \qquad \tilde{E}_{it} = \psi_{eit}\Bigg(\sum_{f\in\mathcal{F}_{it}} \Big(\tilde{\psi}_{fit}\frac{e_{fit}}{\overline{e}_f}\Big)^{\frac{\lambda-1}{\lambda}}\Bigg)^{\frac{\lambda}{\lambda-1}}$$

Relative first-order conditions identify the relative productivity of fuel f as a function of observables up to parameter values:

$$\tilde{\psi}_{fit} = \Big(\frac{p_{fit}}{p_{eit}}\Big)^{\frac{\lambda}{\lambda-1}}\Big(\frac{e_{fit}}{e_{eit}}\Big)^{\frac{1}{\lambda-1}}\frac{\overline{e}_f}{\overline{e}_e} \tag{2.13}$$

The intuition underlying equation (2.13) is that relative fuel productivities equate relative fuel prices to relative marginal products $\frac{p_{fit}}{p_{eit}} = \tilde{\psi}_{fit}^{\frac{\lambda-1}{\lambda}}(\frac{e_{eit}}{e_{fit}})^{\frac{1}{\lambda}}$. I then exploit these optimality conditions by substituting back the implied relative fuel productivity terms (2.13) into the energy production function (2.12) and rearranging:

$$\frac{\tilde{E}_{it}}{\tilde{e}_{eit}} = \psi_{eit}\Bigg(\sum_{f\in\mathcal{F}_{it}} \frac{s_{fit}}{s_{eit}}\Bigg)^{\frac{\lambda}{\lambda-1}} \tag{2.14}$$

Where $s_{fit} \equiv p_{fit}e_{fit}$ is spending on fuel f. The intuition underlying equation 2.14 is fairly straightforward. The left-hand side is the value added of an additional unit of electricity in terms of energy. In contrast, the right-hand side is the contribution of electricity productivity and relative spending on other fuels to that value added. Naturally, higher electricity productivity increases the value added of electricity, and higher spending on other fuels also increases the quantity of energy produced for a given unit of electricity.

The only unobservable left in the energy production function is the productivity of electricity, which is correlated with current period quantities and spending on fuels since it is assumed to be known to plants when choosing fuel quantities. To deal with this issue, I assume that electricity productivity follows an AR(1) Markov process with year and plant-fixed effects.[13] Moreover, I allow for plant-specific fixed effects in the productivity of electricity $\mu_i^{\psi_e}$

$$\ln \psi_{eit} = (1 - \rho_{\psi_e})(\mu_0^{\psi_e} + \mu_r^{\psi_e} + \mu_i^{\psi_e}) + \mu_t^{\psi_e} - \rho_{\psi_e}\mu_{t-1}^{\psi_e} + \rho_{\psi_e} \ln \psi_{eit-1} + \epsilon_{it}^{\psi_e} \tag{2.15}$$

I then take the log of equation (2.14) and use the Markov process above to get an estimating equation:

$$\ln \hat{\tilde{E}}_{it} - \ln \tilde{e}_{eit} = \Gamma_t + \rho_{\psi_e}(\ln \tilde{E}_{it-1} - \ln \tilde{e}_{eit-1}) + \frac{\lambda}{\lambda - 1}\Big(\ln \sum_{f \in \mathcal{F}_{it}} \frac{s_{fit}}{s_{eit}} - \rho_{\psi_e} \ln \sum_{f \in \mathcal{F}_{it-1}} \frac{s_{fit-1}}{s_{eit-1}} \Big) + \mu_i^* + \epsilon_{it}^{\psi_e} \tag{2.16}$$

Where $\Gamma_t = \mu_0^{\psi_e}(1 - \rho_{\psi_e}) + \mu_t^{\psi_e} - \rho_{\psi_e}\mu_{t-1}^{\psi_e}$ is a year fixed-effect and $\mu_i^* = (1 - \rho_{\psi_e})\mu_i^{\psi_e}$ is the normalized plant fixed effect. Since $\epsilon_{it}^{\psi_e}$ is a shock to the productivity of electricity at time t, it is uncorrelated with choices made at time $t - 1$:

$$\mathbb{E}(\epsilon_{it}^{\psi_e} \mid \mathcal{I}_{it-1}) = 0$$

There are two main endogeneity concerns in this model. First, the lagged value added of electricity and the lagged relative spending on other fuels are correlated with the plant fixed effect μ_i^*, which biases the persistence of electricity productivity ρ_{ψ_e}. This is the standard concern in the dynamic panel literature. Second, contemporaneous relative spending on other fuels is correlated with both the fixed effect μ_i^* and the innovation term $\epsilon_{it}^{\psi_e}$ to electricity productivity, which biases the estimate of the elasticity of substitution λ. Blundell and Bond (1998, 2000), and many others show that these concerns can be addressed with properly specified moment conditions. I use the system GMM approach, which combines both level and difference moment conditions as follows:

[13]The choice of these modified AR(1) processes, where the mean is normalized by the persistence, are standard in the dynamic panel literature with short panels (Blundell and Bond, 2023). It ensures that the average of each state variable observed in the data corresponds to the unconditional average of this process. This means that even though the model is estimated from a short panel (between 2 and 8 years), forward simulations multiple years ahead will match the support of the data. It is equivalent to the assumption that the residuals of the productivity distribution follow an AR(1) process rather than electricity productivity itself.

$$\mathbb{E}(\Delta X_{i,t-1}(\mu_i^* + \epsilon_{it}^{\psi_e})) = 0$$
$$\mathbb{E}(X_{i,t-1}\Delta\epsilon_{it}^{\psi_e}) = 0$$

For $X_{i,t-1} \in \{\ln \tilde{E}_{i,t-1} - \ln \tilde{e}_{e,i,t-1}, \ln \sum_{f\in\mathcal{F}_{i,t-1}} \frac{s_{fit-1}}{s_{eit-1}}\}$ and likewise for $\Delta X_{i,t-1}$. Moreover, these moment conditions yield a consistent estimate of the elasticity of substitution λ under the assumption that shocks affecting relative fuel spending are persistent. This assumption is consistent with many geopolitical shocks persistently affecting fuel prices in the market. Lastly, I get standard errors on the elasticity of substitution using the delta method.

2.6 Identification and Estimation of Fixed Fuel Switching Costs

Each plant has access to a set of fuels $\mathcal{F}_{it}$ and is considering all alternative fuel sets for the next period: $\mathcal{F}' \equiv \mathcal{F}_{it+1} \subseteq \mathbb{F} \equiv \{\text{oe,oge,oce,ogce}\}$. Since all state variables s_{it} are assumed to follow a Markovian process, I start from the recursive formulation of the problem. The plant chooses a fuel set next period $\mathcal{F}'$ to maximize the net present value of lifetime profits:

$$V(s_{it}, \epsilon_{it}, \mathcal{F}_{it}) = \max_{\mathcal{F}'\subseteq\mathbb{F}} \Big\{ \pi(s_{it}, \mathcal{F}_{it})/\sigma_\epsilon - \mathcal{K}(\mathcal{F}' \mid \mathcal{F}_{it}, s_{it})/\sigma_\epsilon + \epsilon_{\mathcal{F}'it} + \beta\,\mathbb{E}(V(s_{it+1}, \epsilon_{it+1}, \mathcal{F}') \mid s_{it}) \Big\} \tag{2.17}$$

Where the fuel set switching cost function, $\mathcal{K}(\mathcal{F}' \mid \mathcal{F}_{it}, s_{it})$, was defined in Table 2.5. It is a function of productivity z_{it} and access to natural gas pipelines. σ_ϵ is a parameter that maps units of profits (dollars) to units of the fixed cost shocks.[14] From now on, I define the parameters governing the switching cost function $\theta_1 = \{\kappa_{g1}, \kappa_{g0}, \kappa_c, \gamma_{g1}, \gamma_{g0}, \gamma_c\}$ for coal c and gas g, and θ_2 the parameters underlying the evolution of state variables. I use κ_{g1} to denote the fixed cost of adding natural gas for plants located in a district near the pipeline network and κ_{g0} for plants located in a district that isn't immediately adjacent to the pipeline network, and likewise for salvage values. I make the assumption that cost shocks are iid and come from a standardized Type 1 Extreme value $\epsilon_{\mathcal{F}'it} \sim \textit{Gumbel}(0,1)$. This allows me to analytically integrate these shocks and work with the expected value function, $W(s_{it}, \mathcal{F}_{it}) = \mathbb{E}(V(s_{it}, \epsilon_{it}, \mathcal{F}_{it}))$:

[14] An equivalent approach would be to map units of the fixed cost shocks to units of profits (dollars). Once σ_ϵ is known, I can always switch between dollars and units of the Gumbel distribution.

$$W(s_{it},\mathcal{F}_{it}) = \gamma + \ln\Bigg(\sum_{\mathcal{F}'\in\mathbb{F}} \exp\Big(\underbrace{\pi(s_{it},\mathcal{F}_{it})/\sigma_\epsilon - \mathcal{K}(\mathcal{F}' \mid \mathcal{F}_{it}, s_{it})/\sigma_\epsilon + \beta \int W(s_{it+1},\mathcal{F}')f(s_{it+1} \mid s_{it})ds_{it+1}}_{\upsilon_{\mathcal{F}'}(s_{it},\mathcal{F}_{it})}\Big)\Bigg)$$

Where $\gamma \approx 0.5772$ is the Euler–Mascheroni constant and $\upsilon_{\mathcal{F}'}(s_{it},\mathcal{F}_{it})$ is the choice-specific value function. Then, the probability of choosing fuel $\mathcal{F}'$ has a logit formulation, simplifying the likelihood. Note that this probability is implicitly a function of both θ_1 and θ_2. Next, I discuss the evolution of each state variable.

$$Pr(\mathcal{F}' \mid \mathcal{F}_{it}, s_{it};\theta_1,\theta_2) = \frac{exp\Big(\upsilon_{\mathcal{F}'}(s_{it},\mathcal{F}_{it};\theta_1,\theta_2)\Big)}{\sum_{\mathcal{F}\in\mathbb{F}} \exp\Big(\upsilon_{\mathcal{F}}(s_{it},\mathcal{F}_{it};\theta_1,\theta_2)\Big)}$$

Plants take expectation over all productivity terms, fuel prices, and material prices, which I separate into two categories. Non-selected state variables, which I observe for every plant in every year ($\psi_{oit}, \psi_{eit}, p_{ot}, p_{eit}, z_{it}, p_{mit}$), and selected state variables, which I only observe when plants are using the relevant fuel ($\psi_{cit}, \psi_{git}, p_{cit}, p_{git}$). I assume that plants do not take expectation over the rental rate of capital and aggregate wages to reduce the computational burden of this problem.

2.6.1 Evolution of State Variables

All state variables follow a persistent AR(1) process with time (t) and region (r) fixed effects. To reduce the state space, I assume that the persistence for fuel prices and productivity is the same $\rho_{p_f} = \rho_{\psi_f} = \rho_f$. This assumption allows me to define fuel productivity over prices as a single state variable, which always enters together in plants' profit function p_{fit}/ψ_{fit} through the energy price index. I allow for systematic differences in the productivity/price of each fuel across plants, μ_i^f, which I call *fuel comparative advantage*. $\forall f$,

$$\mathbb{E}\left[\, \ln(\psi_{fit+1}/p_{fit+1}) \mid \mathcal{I}_{it}\right] = (1-\rho_f)(\mu_0^f + \mu_r^f + \mu_t^f + \mu_i^f) + \rho_f \ln(\psi_{fit}/p_{fit}) \tag{2.18}$$

In this context, A positive (log) fuel price shock is isomorphic to a negative (log) fuel productivity shock and vice versa. The region dummies μ_r^f (Northern, North-East, Eastern, Center, Western, Southern) capture the bulk of the spatial variation in fuel prices and technology. I also assume a similar process for hicks-neutral productivity $\ln z_{it}$ and for the (log) price of materials $\ln p_{mit}$. These two state variables and the price/productivity of oil and electricity were recovered in previous sections for all plants in all years. As such, these Markovian processes can be estimated directly using the system GMM approach of Blundell and Bond (2000). Following Bonhomme and Manresa (2015), I reduce the dimension of oil and electricity comparative advantages μ_i^f by grouping plants into clusters of unobserved heterogeneity using K-means clustering.

However, not every plant uses gas and coal. When plants are not using gas or coal in year t, the price/productivity process starts at the initial condition (equivalent to $t = 0$), and plants take expectation over idiosyncratic shocks. $\forall f \in \{coal, gas\}$:

$$\mathbb{E}\left[\ln(\psi_{fit+1}/p_{fit+1}) \mid f \notin \mathcal{F}_{it}\right] = \mu_0^f + \mu_t^f + \mu_r^f + \mu_i^f$$

Estimating the distribution of plant-specific comparative advantages for coal and gas only from plants that currently use gas or coal is likely to be biased due to selection. Recovering the distribution of μ_i^f from selected plants may not reflect the distribution across all plants, which would bias fixed cost estimates. In the next section, I show how to recover the distribution of these fuel comparative advantages for plants that are not currently using gas or coal jointly with fixed costs following the approach of Arcidiacono and Jones (2003). Using this approach, I recover the distribution of fuel comparative advantages that is most likely to rationalize observed fuel set choices under the assumption that plants know their comparative advantage and use that information to make decisions. Lastly, I allow all shocks to state variables to be arbitrarily correlated in a multivariate normal distribution with mean zero and a positive semi-definite covariance matrix $\mathbf{\Sigma}$. I discretize the entire state space following Farmer and Akira Toda (2017).

$$(\epsilon_{it}^o, \epsilon_{it}^e, \epsilon_{it}^g, \epsilon_{it}^c, \epsilon_{zit}, \epsilon_{mit}) \equiv \boldsymbol{\epsilon_{it}} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$$

2.6.2 Identification of Fixed Costs and Fuel Comparative Advantages

To learn about the extent to which the distribution of comparative advantage for natural gas and coal is selected, I follow the algorithm proposed by Arcidiacono and Jones (2003). I assume that the distribution of comparative advantages comes from finite mixtures with $K = 3$ points of support for each fuel. I parameterize the initial guess of the mean and variance of the finite mixture to the mean and variance of the empirical (selected) distribution ($\tilde{\mu}_f, \tilde{\sigma}^2_{\mu_f}$):

$$\sum_k^K \pi^0_{fk}\mu_{fk} = \tilde{\mu}_f \qquad\qquad \sum_k^K (\mu_{fk} - \tilde{\mu}_f)^2 \pi^0_{fk} = \tilde{\sigma}^2_{\mu_f}$$

Where $\pi^0_{fk} = Pr(\mu_{fk})$ is the unconditional probability of being type k, where types refer to support points of the fuel comparative advantage distribution, and $\sum_k \pi^0_{fk} = 1$. In this context, external estimation of parameters governing the distribution of random effects from a selected sample of plants that use these fuels leads to biased estimates of $\tilde{\mu}_g, \tilde{\mu}_c, \tilde{\sigma}^2_{\mu_g}, \tilde{\sigma}^2_{\mu_c}$. Indeed, plants with a larger comparative advantage for coal are more likely to use coal, and likewise for gas. Thus, I expect to get upward biases in both the mean of coal and gas. Using the law of total probability, I can integrate the unconditional distribution of comparative advantages using the full information (log) likelihood. Assuming there is only one finite mixture over both coal and gas for notation convenience, and where the distribution of comparative advantages are independent across fuels such that $\pi_k \in \Pi = vec(\Pi_g \otimes \Pi_c)$, where $\pi_{kg} \in \Pi_g$ and $\pi_{kc} \in \Pi_c$:

$$\ln \mathcal{L}(\mathcal{F}, s \mid \theta_1, \theta_2) = \sum_{i=1}^n \ln\Bigg[\sum_k \pi_k \Bigg[\prod_{t=1}^T Pr(\mathcal{F}_{it+1} \mid \mathcal{F}_{it}, s_{it}, \mu_i = \mu_k; \theta_1, \theta_2)\Bigg]\Bigg] + \sum_{i=1}^n \sum_{t=1}^T \ln f(s_{it} \mid s_{it-1}; \theta_2) \tag{2.19}$$

In particular, the likelihood in (2.19) assumes that the state transitions are independent of the distribution of comparative advantages for coal and gas.[15] This is possible if the parameter estimates $\hat{\theta}_2$ are unbiased from selected data. Initially, the true probability weights π_k over the support of the finite mixture are unknown due to selection, but Arcidiacono and Jones (2003), Arcidiacono and Miller (2011) provides a method to recover the unselected distribution by sequentially iterating over the fixed costs to maximize the likelihood and updating the probability weights $\pi^0_k, \pi^1_k, \pi^2_k, \ldots$ using an EM algorithm. Following Bayes' law, one can show that the solution to this maximum likelihood problem is the same as the solution to a sequential EM algorithm that uses the posterior conditional probabilities that plant i is of type k given all observables, including choices made:

[15]This assumption isn't necessary, but it simplifies the computation of the model in the presence of these comparative advantages.

$$
\begin{aligned}
\hat{\theta}_1 &= \underset{\theta_1,\theta_2,\pi}{\arg\max} \sum_{i=1}^{n} \ln\Bigg[\sum_k \pi_k \Bigg[\prod_{t=1}^{T} Pr(\mathcal{F}_{it+1} \mid \mathcal{F}_{it}, s_{it}, \mu_i = \mu_k; \theta_1, \theta_2) \Bigg]\Bigg] \\
&\equiv \underset{\theta_1}{\arg\max} \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_k \rho(\mu_k \mid \mathcal{F}_i, s_i; \hat{\theta}_1, \hat{\theta}_2, \hat{\pi}) \ln Pr(\mathcal{F}_{it+1} \mid \mathcal{F}_{it}, s_{it}, \mu_i = \mu_k; \theta_1, \hat{\theta}_2)
\end{aligned}
$$

Where $\mathcal{F}_i$ is the sequence of fuel set choices I observe establishment i making. Using Bayes' rule, the conditional probability that plant i is of type k is given by the current guess of the unconditional probability $\hat{\pi}_k$ weighted by the probability that the plant makes the observed sequence of fuel set choices conditional being type k:

$$
\rho(\mu_k \mid \mathcal{F}_i, s_i; \theta_1, \theta_2, \hat{\pi}) = \frac{\hat{\pi}_k \Bigg[\prod_{t=1}^{T} \Bigg[\prod_{\mathcal{F} \subseteq \mathbb{F}} \Big[Pr(\mathcal{F}_{it} \mid s_{it}, \mu_i = \mu_k; \theta_1, \theta_2) \Big]^{\mathbb{I}(\mathcal{F}_{it} = \mathcal{F})} \Bigg]\Bigg]}{\sum_k \hat{\pi}_k \Bigg[\prod_{t=1}^{T} \Bigg[\prod_{\mathcal{F} \subseteq \mathbb{F}} \Big[Pr(\mathcal{F}_{it} \mid s_{it}, \mu_i = \mu_k; \theta_1, \theta_2) \Big]^{\mathbb{I}(\mathcal{F}_{it} = \mathcal{F})} \Bigg]\Bigg]} \tag{2.20}
$$

The idea underlying the EM algorithm is to iteratively estimate fixed cost parameters θ_1 given some guess of the distribution of comparative advantages $\{\pi_k\}_k$ – M step, draw new comparative advantages using Baye's law from (2.20), which are used to update the unconditional distribution of comparative – E step, and repeat this procedure until the likelihood in (2.19) is minimized. Details of the algorithm can be found in Appendix C.3.

2.7 Estimation Results – Steel Manufacturing

Outer Production Function Estimation Results

Estimates of the outer production function parameters can be found in Table 2.6. The average output and revenue elasticities with respect to intermediate materials are much larger than those with respect to other inputs and are consistent with the literature (Doraszelski and Jaumandreu, 2013, 2018, Gandhi et al., 2020, Grieco et al., 2016). This is primarily due to the importance of iron ore in steel production. Average output and revenue elasticities are considerably larger for energy than labor and capital due to the large quantities of fuels required to produce steel. The estimated demand elasticity is also consistent with estimates by Zhang (2019), who finds a demand elasticity of around 4 in the Chinese Steel industry. Using these estimates, I can construct estimates of the price $p_{\hat{E}_{it}}$ and quantity of the energy bundle for each plant $\hat{E}_{it}$ from the relation first-order conditions in equation 2.8, which I use to estimate the energy production function.

Production and Demand Parameters		Average Output Elasticities		Average Revenue Elasticities
Elasticity of substitution $\hat{\sigma}$	1.80 [1.374,3.054]	Labor	0.040 [0.037,0.048]	0.030 [0.029,0.030]
Returns to scale $\hat{\eta}$	1.23 [1.137,1.444]	Capital	0.023 [0.014,0.035]	0.017 [0.010,0.025]
Elasticity of demand $\hat{\rho}$	3.84 [2.695,4.914]	Materials	1.008 [0.935,1.185]	0.745 [0.741,0.749]
Elasticity of outside good $\hat{\theta}$	0.63 [0.552,0.683]	Energy	0.155 [0.144,0.182]	0.115 [0.113,0.117]
Observations	8,547			

Bootstrap 95% confidence interval in brackets (499 reps)

Table 2.6: Estimates of Outer Production Function

Notes: The average output (revenue) elasticities are defined as the average of the individual output (revenue) elasticity, where the output elasticity is $\frac{\partial y_{it}}{\partial x_{jit}} \frac{x_{jit}}{y_{it}}$ for $y_{it} \in \{Y_{it}, R_{it}\}$ and $x_{jit} \in \{L_{it}, K_{it}, M_{it}, E_{it}\}$

Energy production function estimation results

	Steel
Elasticity of substitution $\hat{\lambda}$	2.224^{***}
	(0.231)
Persistence of electricity productivity $\hat{\rho}_{\psi_e}$	0.649^{***}
	(0.118)
Observations	3,459

Standard errors in parentheses

$^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.7: Estimates of Energy Production Function

Notes: I use the delta method to recover the standard error of $\hat{\lambda}$ where $\hat{\sigma}_{\lambda} = (\hat{\lambda} - 1)\hat{\sigma}_{\gamma}$. Moreover, the number of observations in the energy production function (3,459) is lower than in the outer production function (8,547). This is because the method to estimate the energy production function constructs moments that require at least 3 years of observation per plant to yield consistent estimates (Blundell and Bond, 1998, 2000).

Turning to the energy production function, results indicate that the elasticity of substitution between fuels $\hat{\lambda}$ is larger than the elasticity of substitution between energy and non-energy inputs $\hat{\sigma}$ from Table 2.6. This is important because the larger the elasticity of substitution between fuels, the larger the aggregate gains from carbon taxation (Acemoglu, Aghion, Bursztyn and Hemous, 2012). More substitution possibilities mean more emission reduction can be achieved by substituting away from polluting fuels rather than by reducing output, which is a key trade-off in evaluating carbon policy. Next, I construct estimates of the fuel-specific productivity for each plant in each year $\hat{\psi}_{fit}$ and discuss its implications.

Heterogeneity in fuel productivity across fuel sets

Results are summarized in Figure 2.9. There are a few takeaways. First, electricity is by far the most productive fuel, averaging between two and three times the productivity of other fuels. When comparing coal and natural gas, one mmBtu of coal is, on average, 30% more productive than one mmBtu of gas. However, this gap expands to 185% when looking at productivity per dollar because coal is significantly cheaper – averaging one-fifth of the price of natural gas. This can explain the prevalence of coal in the Indian steel industry.

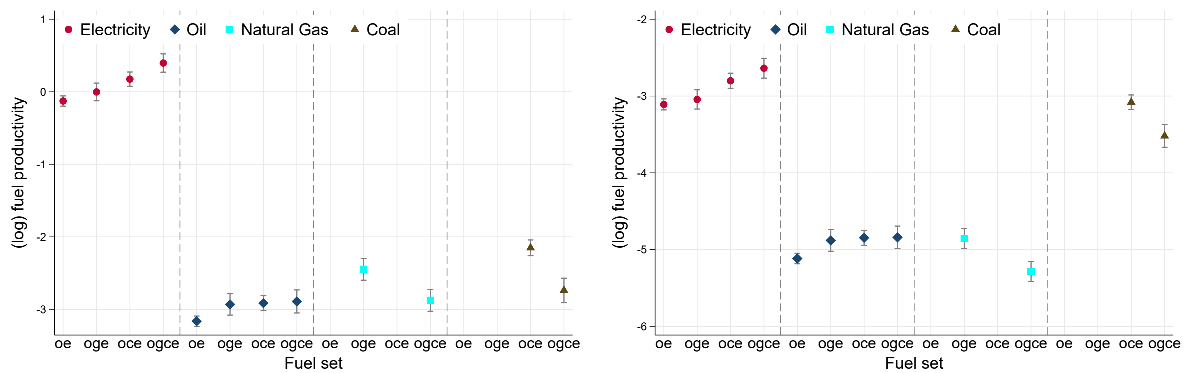

(a) Fuel productivity per mmBtu – $\ln(\psi_{fit}/\overline{e}_f)$ (b) Fuel productivity per dollar – $\ln(\psi_{fit}/(\overline{e}_f * p_{fit}))$

Figure 2.7: Mean (Log) Fuel Productivity Among Plants Using Different Fuel Sets

Notes: The figure is created by taking the sample average of the estimated log productivity for all four fuels by fuel sets. Fuel set labels are created as follows: oe = oil and electricity, oge = oil, gas, and electricity. oce = oil, coal, and electricity. ogce = oil, gas, coal, and electricity. I divide by the geometric mean of fuel quantities $\overline{e}_f$ because fuel productivity is originally in normalized units due to the normalization in estimation.

Estimation of Fixed Costs

Fixed cost estimates are reported in Table 2.8. These estimates encompass the tangible expenses related to new fuel-burning technologies and intangible costs associated with fuel adoption. This includes logistical challenges, new contractual agreements for transportation and storage, as well as potential opportunity costs from diverting labor away from production. These costs are substantial, ranging from 28 to 40 million dollars, and align well with the upper echelon of existing accounting estimates.[16]

		Fixed Costs (Million USD)	Salvage Values (Million USD)
Natural Gas	*Pipeline Access*	28.83	15.12
	No Pipeline Access	40.46	
Coal		28.82	8.33
Total Factor Productivity (100 % increase)		0.82	0.25
Observations		2,393	

Table 2.8: Estimates of Fuel Set Fixed Costs and Salvage Values

Notes: This table shows the fixed cost and salvage value estimates for each fuel in million U.S. dollars. For natural gas, these costs vary based on whether plants are in a district with access to a natural gas pipeline. The parameter in front of "Total Factor Productivity" is the effect of doubling productivity on the fixed costs and salvage values of any fuel and is meant to capture how these costs vary with plant size. The sample size is lower than the energy production function because I removed the last year of observation since I don't observe subsequent fuel set choices.

Coal adoption tends to be 30% cheaper than gas adoption. Moreover, plants without access to high-pressure natural gas pipelines incur 40% higher adoption costs due to the need for alternative transportation methods, such as liquefied natural gas (LNG), which can be costly. This effect of pipeline accessibility is consistent with findings from Scott (2021) in his study of U.S. power plants. The observed salvage values for coal and natural gas are significantly lower, ranging from 57% to 71% below the fixed costs. While fixed costs are nominally very large, the role of plant size is relatively small, as raising productivity by 1% only leads to an $\$8,200$ increase in fixed costs and a $\$2,500$ increase in salvage values. Importantly, the combination of substantial fixed costs and relatively low salvage values likely contributes to situations of technology lock-in, which I discuss in the next section.

2.7.1 Selection Bias in Fuel Productivity – Evidence of Technology Lock-in

The problem of technology lock-in is pervasive, as the Indian Ministry of Steel reports that inefficient plants face difficulties in transitioning out of old technologies:

> "*The higher rate of energy consumption is mainly due to obsolete technologies*

[16]While recent comprehensive accounting estimates of switching costs are hard to find, a single electric arc furnace may cost between a few hundred thousand dollars and a few million dollars (Source: Alibaba's listings `https://www.alibaba.com/product-detail/WONDERY-Custom-Made-Siemens-PLC-Industrial_1600732474634.html`), whereas switching from pig iron, typically produced with a coal-powered blast furnace, to direct reduced iron, typically produced with gas-powered oxygen furnaces would historically cost upwards of USD 70 millions Miller (1976).

> *including problems in retrofitting modern technologies in old plants, old shop floor & operating practices*" Indian Ministry of Steel (2023)

To understand factors that prevent this transition, I revisit the distribution of fuel productivity by considering selection bias in the distribution of fuel comparative advantages. I find significant evidence of selection bias for both coal and natural gas. Indeed, plants that do not use natural gas would be 30% less productive at using natural gas than plants that do, whereas this effect goes up to 80% for coal. Combined with high fixed costs, this productivity gap undermines switching from coal to natural gas and exacerbates technology lock-in. This is because plants that do not currently use natural gas have less to gain from paying the fixed costs, whereas plants that currently use coal have little to gain from dropping coal.

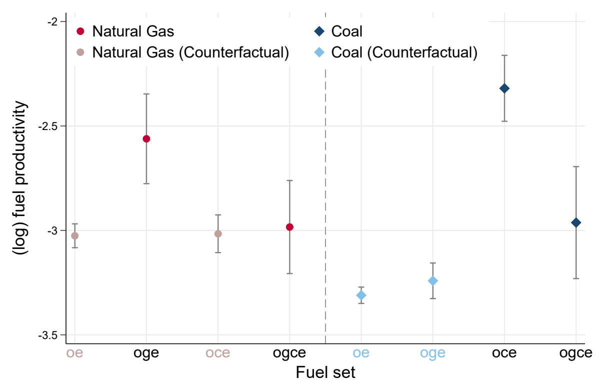

Figure 2.8: Distribution of Fuel Productivity – Including Counterfactual Fuel Sets

Notes: This figure shows the distribution of fuel productivity per mmBtu ($\ln \psi_{fit}/\overline{e}_f$) with 95% confidence intervals for coal and natural gas and includes counterfactual productivity for plants with fuel sets that exclude gas and/or coal. The distribution of fuel productivity for counterfactual fuel sets was computed by simulating draws from the estimated distribution of unobserved heterogeneity (*comparative advantages*) in the dynamic discrete choice model, using the conditional probability distribution $\rho(\mu_k \mid \mathcal{F}_i, s_i; \hat{\theta}_1, \hat{\theta}_2, \hat{\pi})$.

I further strengthen this argument by showing that plants with more fuel tend to face a lower marginal cost of energy in Table 2.9. This difference is largely explained by existing variation in fuel productivity rather than the additional substitution margin that a new fuel provides. In such a context, plants have little incentive to pay the large fixed costs required to add natural gas.

Spatial differences also explain much of the technological lock-in in the industry. As discussed previously with Figure 2.1, many plants using large coal-powered blast furnaces are located in the "Steel Belt" near coal and iron ore mines, as part of Eastern India. At the same time, the natural gas pipeline network is developed in Western India but undedeveloped in Eastern India. While Table 2.8 showed that the direct cost of natural gas increases by 50% without pipeline

		OCE	OGE	OGCE
Total Difference	Percent (%) Difference with OE	-65.65	-71.54	-86.97
Option Value		36.14	5.42	6.3
Fuel Productivity	Percent (%) of Total	62.6	97.75	94.84
Fuel Prices		1.25	-3.18	-1.14

Table 2.9: Shapley Decomposition — Difference in Average Marginal Cost of Energy Between Fuel Sets

Notes: I compare the observed differences in the average (across plants) marginal cost of realized energy between plants who use coal and/or gas on top of oil and electricity (oce, oge, ogce) relative to plants who only use oil and electricity (oe).

access, Figure 2.9 suggests that the opportunity cost for large coal users in Eastern India would also be very large due to their high comparative advantage at using coal relative to other fuels.

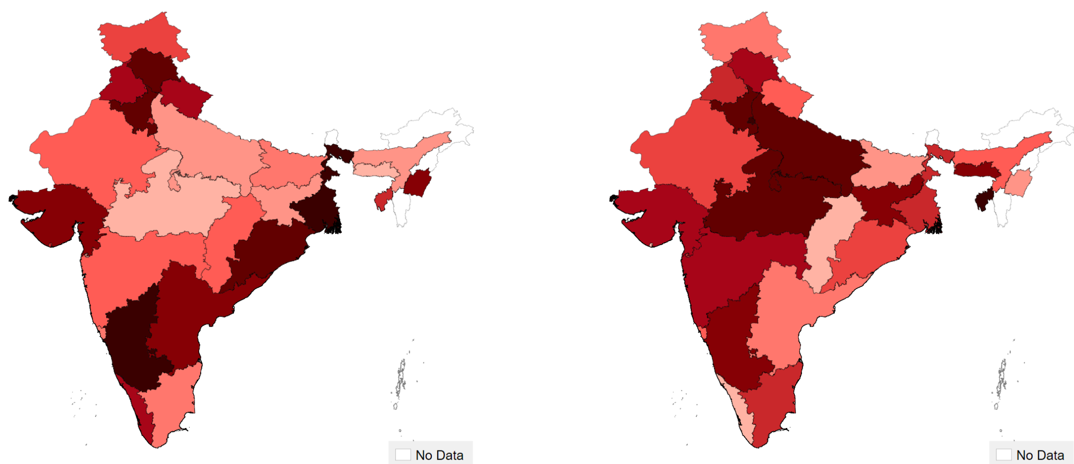

(a) Coal (b) Natural Gas

Figure 2.9: Spatial Distribution of Estimated Fuel Productivity

Notes: The figures plot the distribution of productivity for coal and natural gas across Indian States. Each shade corresponds to a quantile. Darker shades of red are regions where plants are more productive at using the given fuel.

Model fit

Overall, the estimates of switching costs allow the model to predict quite well the empirical distribution of fuel set choices and the observed transition patterns between fuel sets. The model does slightly worse at predicting the transitions for plants that start with all four fuels because it only represents 8% of the sample. The blue bars (model) are constructed in all figures below by adding the predicted probability that each plant uses each fuel set, integrated over the conditional distribution of comparative advantages.

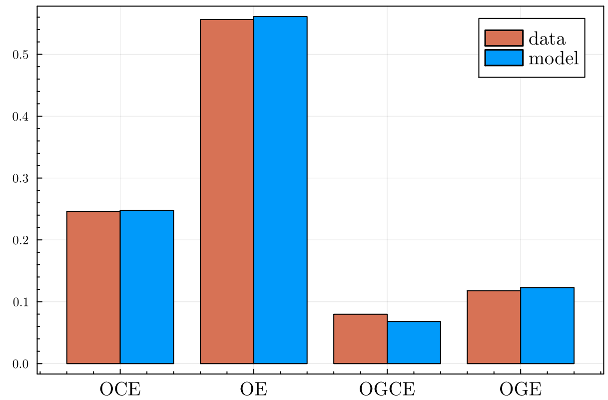

Figure 2.10: Unconditional Distribution of Fuel Sets, Model vs. Data ($N = 2,393$)

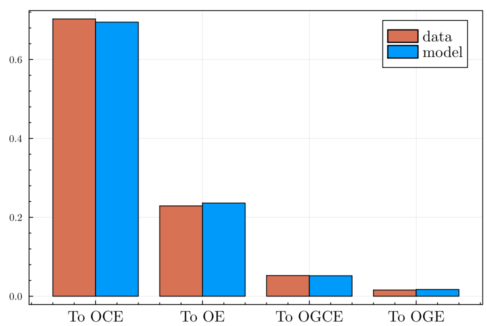

From Oil, Coal, Electricity (OCE) – $N = 572$

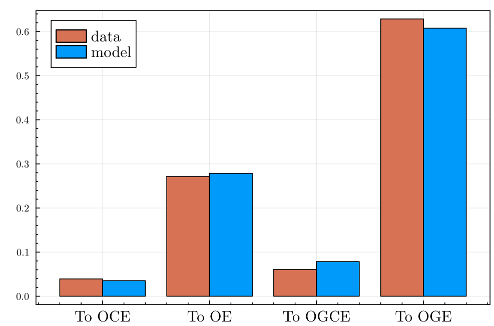

From Oil, Gas, Electricity (OGE) – $N = 280$

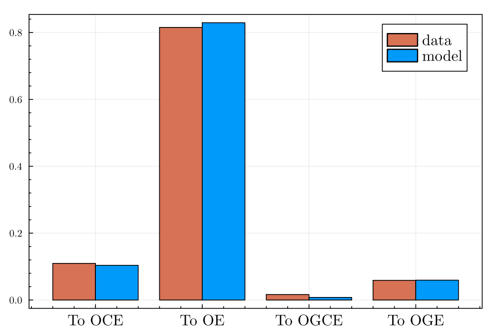

From Oil, Electricity (OE) – $N = 1,342$

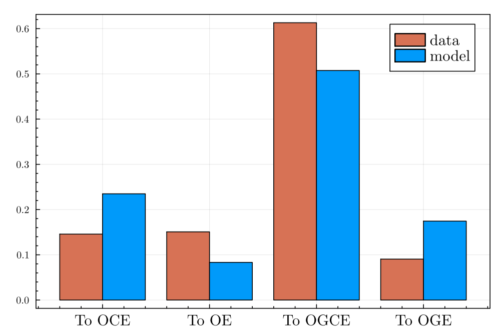

From Oil, Gas, Coal, Electricity (OGCE) – $N = 199$

Figure 2.11: Conditional Distribution of Fuel Sets (Transition), Model vs. Data

Furthermore, I look at the relationship between plant size (proxied by Hicks-neutral productivity) and fuel switching and find a positive relationship between the probability that a plant adds a new fuel at $t + 1$ and its productivity at t, both in the data and the model. See Figure 2.12.

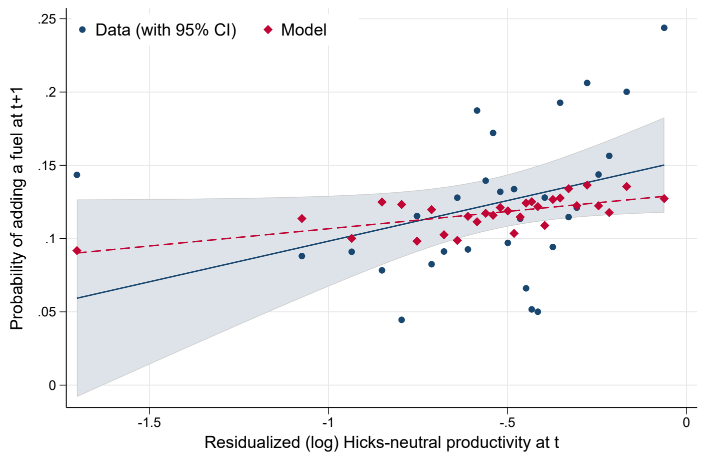

Figure 2.12: Adding a Fuel at $T + 1$ Against Hicks-Neutral Productivity at T

This figure is a binned scatter plot projecting fuel switching $t + 1$ against Hicks-neutral productivity at t. For the data, I make a linear projection of whether plants added a fuel between t and $t + 1$ against productivity. I do the same for the model, replacing choices made by the predicted choice probability. I control for both year fixed effects and input prices in the model and in the data.

2.8 Externality Mitigation Policies

In this section, I study the effectiveness of various policies in mitigating externality damages from fuel combustion to improve social welfare. I detail how externality damages are constructed and perform two counterfactual policy experiments. First, I quantify the trade-off between emission reduction and output for various levels of fossil fuel taxes, where the tax rate is proportional to each fuel's emission intensity (carbon tax). I compare this trade-off with an economy without heterogeneity in fuel productivity. Second, I discuss the pervasiveness of technology lock-in in this economy and evaluate a potential solution that uses proceeds from the carbon tax to finance a subsidy that reduces the fixed cost of natural gas adoption.

Externality Damages

Externality comes from the release of pollutants in the air by the combustion of fuels. All pollutants are converted into carbon dioxide equivalent (CO_{2e}) using standard scientific calculations from the U.S. EPA. Then, each unit of fuel f's potential energy contributes to contemporaneous greenhouse gas emissions as follows: 1 *mmBtu* of e_f releases γ_f short tons of CO_{2e}. γ_f are fuel-specific emission intensities calculated using the global warming potential (GWP) method detailed in Appendix A.2. For example, 1 *mmBtu* of coal releases roughly twice as much carbon dioxide equivalent in the air as 1 *mmBtu* of natural gas $\frac{\gamma_c}{\gamma_g} \approx 2$. Fuel-specific emission intensities define the relative tax rate between different fuels. See table 2.10.

Fuel	Average Price (Rupees)	Emission Factor (Kg CO_{2e}/mmBtu)
Coal	262	100
Oil	665	82
Electricity	1,681	65
Natural Gas	1,307	60

Table 2.10: Example of Average Fuel Prices With and Without Carbon Tax

Notes: These prices are averaged across all sample years. At the baseline, they reflect the average prices that plants pay for each fuel within a year. Natural gas is slightly less polluting than electricity. This is because the vast majority of Indian electricity is generated with either coal or renewables such as hydro.

Since the carbon tax is a per-unit tax on fossil fuels $p_{fit} + \tau_f$, I separate the evolution of fuel prices/productivity of Equation 2.18 into two separate processes: one for prices and one for productivity separately. In practice, from the discrete grid that forms the Markov chain describing the evolution of fuel price/productivity, I construct two separate grids, one for prices and another for productivity. Grid points are found by matching moments of the discrete Markov chain for fuel prices and productivity with moments from observed fuel prices in the data and estimated fuel productivity. Details can be found in Appendix D.1. Intuitively, this is akin to decomposing the shocks to fuel prices/productivity into a shock to fuel prices and a fuel productivity shock.

2.8.1 The Equimarginal Principle With Multidimensional Heterogeneity

While both multidimensional heterogeneity in fuel productivity and dynamic fuel switching complicate the graphical analysis of the equilibrium under a carbon tax, the equimarginal principle still holds. A carbon tax levied on fossil fuels not only equalizes marginal abatement costs (expressed in forgone profit) across plants, it also equalizes marginal abatement costs across fuels within a plant. This can be seen by rearranging first-order conditions of the plant's static profit maximization, where the carbon tax τ is levied on firms' total emissions $CO_{2e,it} = \sum_{f \in \mathcal{F}_{it}} \gamma_f e_{fit}$:

$$\left[\frac{\partial Y_{it}}{\partial e_{fit}}\Big(P(Y_{it}) + Y_{it}P'(Y_{it})\Big) - p_{fit}\right] \Big/ \gamma_f = \tau \quad \forall f \in \mathcal{F}_{it} \tag{2.21}$$

The left-hand side is plant i's marginal cost of abating one ton of CO_{2e} using fuel f (in terms of forgone profits), whereas the right-hand side is the carbon tax rate. This equation is key because it provides a microfoundation for marginal abatement costs. In this highly heterogeneous context, firms have multiple ways to equalize their marginal abatement costs. In the short run, they can vary the quantity of each fuel in their set or their output. In the long run, they can pay fixed switching costs to adopt new fuels or drop existing ones. Critically, because fuel marginal products $\frac{\partial Y_{it}}{\partial e_{fit}}$ vary across fuels and across plants, plants will make different decisions to equalize their marginal abatement costs. This is the benefit of a market-based policy. Despite facing the same tax schedule, plants may be more or less targeted by the tax based on input prices they face, their fuel sets, how productive they are overall, and how productive they are at using different fuels.

2.8.2 Carbon Tax and the Trade-off Between Output and Emission Reduction

In Figure 2.13, I trace the trade-off between output and emission reduction for various carbon tax rates. Each point on the curve corresponds to a different level of the carbon tax, and together, they form a production frontier in output and emission reduction. I simulate the economy with and without the carbon tax for 40 years starting from 2016 and look at the net present value (NPV) of outcomes along the entire path. For $X = \{\text{aggregate output, aggregate emissions}\}$ and a given carbon tax rate τ:

$$\mathbb{E}(X(\tau)) \approx \frac{1}{S}\sum_{s=1}^{S}\sum_{t=0}^{40}\beta^t X_{ts}(\tau)$$

As the level of the tax approaches zero, the model converges to the no-tax economy with 100% of output and 0% of emission reduction. As the tax level increases, emissions decrease but so does aggregate output. The production frontier is concave because of the increasing marginal cost of reducing aggregate emissions, consistent with previous findings by Fowlie et al. (2016). Fuel substitution (and more generally input substitution) is more effective initially, where much emission reduction can be achieved by substituting coal with cleaner fuels such as natural gas and electricity. However, as the carbon tax rate increases, more emission reduction comes at the cost of plants scaling down their operation, decreasing aggregate output because marginal plants have already substituted cleaner fuels for dirtier fuels.

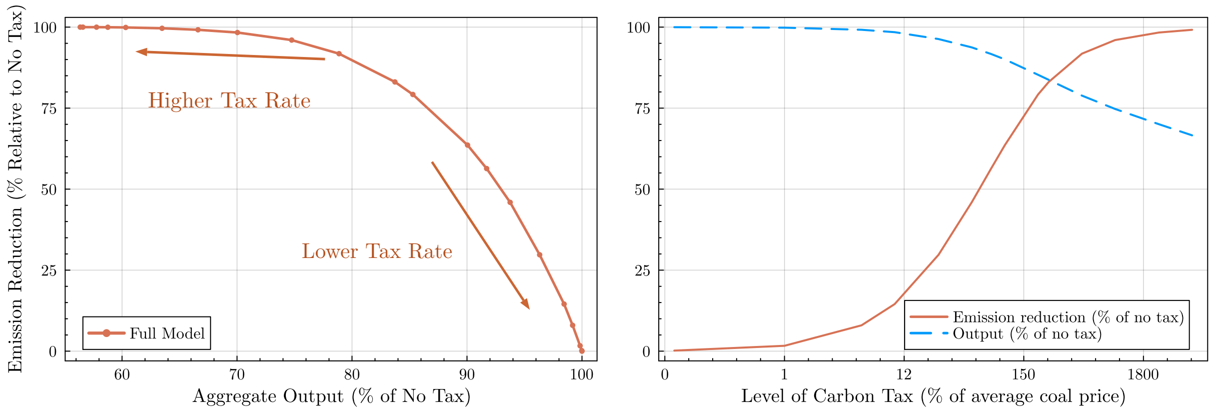

(a) Trade-off between Output and Emission Reduction

(b) Emission Reduction and Output Across Tax Levels

Figure 2.13: Production Frontier in Output and Emission Reduction for Various Carbon Tax Rates

Notes: This production frontier was constructed by simulating the economy under 21 different carbon tax levels, ranging from 0 (no tax) to approximately infinity. Linear interpolation is assumed for the trade-off between each tax level. As the level of the tax approaches infinity, the aggregate output does not reach 0. This is a feature of the CES production function. Indeed, as fuel prices are extremely high, fuel consumption approaches zero, but plants always use some positive amount of fuel.

The Role of Heterogeneity in Fuel Productivity – Intensive Margin

I compare in Figure 2.14 what happens when removing heterogeneity in fuel productivity from the economy to highlight its role in the trade-off. To do so, I re-estimate an energy production function in which plants have heterogeneous energy productivity ψ_{Eit} but have the same average fuel productivity. Details on the estimation of this production function are in Appendix D.3, and follow the dynamic panel approach, similar to the energy production function with heterogeneity in fuel productivity. Estimating this production function matches average fuel quantities and aggregate emissions levels but misses the heterogeneity in fuel shares across plants.

$$E_{it} = \psi_{Eit}\left(\sum_{f\in\mathcal{F}_{it}} \beta_f e_{fit}^{\frac{\lambda-1}{\lambda}}\right)^{\frac{\lambda}{\lambda-1}}$$

The blue production frontier in figure 2.14 corresponds to this economy and shrinks inwards compared to the baseline economy. The restricted economy operates at 88% of no-tax output when reducing emissions by 50%, with an implied elasticity between emission reduction and output of 4.17 – compared to 93.5% of no-tax output in the flexible economy, with an implied elasticity of 7.7. Allowing for heterogeneity in fuel productivity diminishes how much output

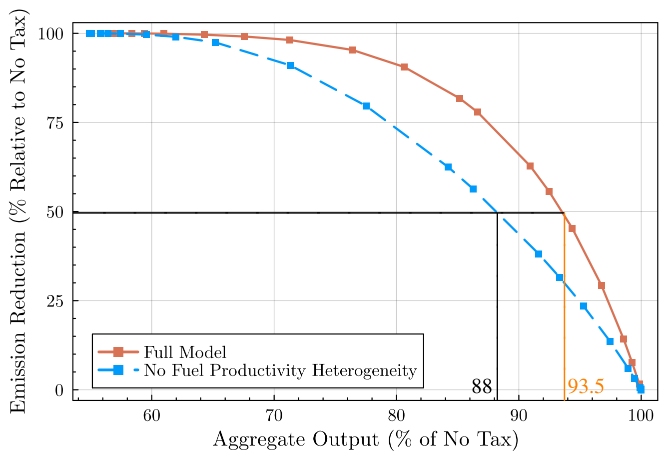

Figure 2.14: Comparison of Trade-off Across Model Specification

Notes: This production frontier highlights the role of the intensive margin only. It compares the trade-off across contemporaneous taxes, excluding forward simulations and switching between fuel sets (extensive margin). The comparison of trade-offs is the same when allowing for the extensive margin to adjust in a forward simulation because heterogeneity in fuel productivity affects contemporaneous reallocation across plants.

must decrease to achieve any reduction in emissions because it increases the aggregate elasticity of substitution between fuels by reallocating output from dirty to clean plants. Two channels explain this reallocation.

First, conditional on fuel prices and fuel sets, the elasticity of the marginal cost of energy with respect to relative fuel prices is not constant across plants. For example, as the price of coal increases relative to the price of gas, plants that are more productive at using coal relative to gas face a larger percentage increase in their cost of energy. This is because larger coal productivity induces specialization in coal, making them more exposed to the relative price increase, as long as fuels are gross substitutes $\lambda > 1$. Since the carbon tax effectively increases the relative price of polluting fuels, plants that are more productive at using polluting fuels are especially hurt by the carbon tax. To see this, let $\tilde{p}_{cit} = p_{cit}/p_{git}$ be the price of coal relative to gas and likewise for relative fuel productivity $\tilde{\psi}_{cit} = \psi_{cit}/\psi_{git}$. Then,

$$\frac{\partial \ln p_{E_{it}}}{\partial \ln \tilde{p}_{cit}} = \frac{\left(\tilde{p}_{cit}/\tilde{\psi}_{cit}\right)^{1-\lambda}}{\sum_{f\in\mathcal{F}_{it}}\left(\tilde{p}_{fit}/\tilde{\psi}_{fit}\right)^{1-\lambda}} = \frac{p_{cit}e_{cit}}{\sum_{f\in\mathcal{F}_{it}} p_{fit}e_{fit}} \tag{2.22}$$

Under cost-minimization, the elasticity of the marginal cost of energy with respect to the relative price of any fuel (e.g., coal relative to gas) is just the plant-specific spending share of that fuel relative to all fuels. This is an application of the Envelope Theorem. Details in Appendix

D.3. Most importantly, this elasticity is increasing in relative fuel productivity. This means that conditional on fuel prices and fuel set, plants that are more productive at using coal spend more on coal and are more sensitive to relative changes in the price of coal:

$$\frac{\partial^2 \ln p_{E_{it}}}{\partial \ln \tilde{p}_{cit} \partial \tilde{\psi}_{cit}} = \frac{(\lambda-1)\psi_{cit}^{\lambda-2}\tilde{p}_{cit}^{1-\lambda}\Big[\sum_{f\in\mathcal{F}_{it}\setminus c}(\tilde{p}_{fit}/\tilde{\psi}_{fit})\Big]}{\big(\sum_{f\in\mathcal{F}_{it}}(\tilde{p}_{fit}/\tilde{\psi}_{fit})^{1-\lambda}\big)^2} > 0 \qquad \text{if} \quad \lambda > 1$$

In contrast, in the economy without fuel-specific productivity, the elasticity of the price of energy with respect to relative fuel prices is constant up to fuel prices and fuel sets.

Second, this heterogeneous increase in marginal energy costs makes polluting plants less competitive than cleaner plants, consistent with the aggregation result by Oberfield and Raval (2021). The tax thus induces a reallocation of output from more polluting to less polluting plants, which increases aggregate fuel substitution and diminishes how much aggregate output must be reduced to achieve any emission reduction target[17]. Moreover, this reallocation channel is a function of the elasticity of demand and returns to scale. Indeed, as the elasticity of demand increases and different output varieties become more substitutable with each other, any variation in relative marginal costs across plants will lead to a larger reallocation of output. In Appendix D.5, I confirm this intuition by showing that the difference between both production frontiers expands as the elasticity of demand increases. In summary, not allowing for this rich heterogeneity between fuel productivity shuts down the reallocation of output from high-emission to low-emission plants and decreases the effectiveness of a carbon tax.

This result can be summarized in Figure 2.15, where I decompose emissions reduction along different levels of the carbon tax. Following the decomposition method proposed by Levinson (2015), I decompose emissions reduction into a technique, composition, and scale effect. The technique effect corresponds to plant-level fuel substitution, whereas the composition effect corresponds to output reallocation across firms, and the scale effect corresponds to variation in aggregate output.

Relative to a model in which all fuels are equally productive, I find that a much larger portion of emissions reduction comes from output reallocation and much less from plant-level fuel substitution when allowing for heterogeneity in fuel productivity. For a carbon tax of 40 USD

[17]Note that the correlation between fuel productivity and total factor productivity (TFP) also matters for this result. Indeed, if plants more affected by the carbon tax were also initially more productive overall, this reallocation effect may reduce aggregate TFP and aggregate output. However, I show in Appendix D.4 that the opposite is true. Plants with higher fuel productivity tend to be less productive overall.

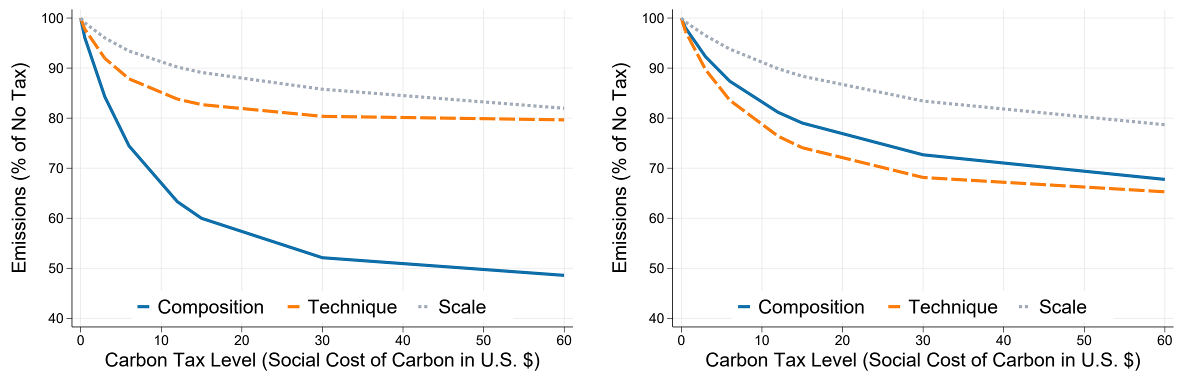

(a) Baseline (b) Equally Productive Fuels

Figure 2.15: Decomposition of Emissions Reduction Under Different Levels of Carbon Tax

Notes: This figure decomposes emissions reduction into three channels following the decomposition method of Levinson (2015) for different levels of the carbon tax. While the level of the carbon tax is expressed in U.S. dollars, they do not account for purchase power parity (PPP) differences between India and the United States. For example, the 40 USD carbon tax is closer to 150 when accounting for PPP differences.

per ton of CO_{2e} (equivalent to 150 USD when adjusted for purchasing power parity). Heterogeneity in fuel productivity induces plants to specialize in different fuels, which then face a higher opportunity cost for switching away from their most productive fuels. As a result, this leads to technology lock-in (less fuel switching) and more variation in exposure to the tax (more reallocation of output).

To benchmark this result with the literature, I do two exercises. First, I compare the aggregate trade-off between output and emissions with Fowlie et al. (2016), who conduct similar policy exercises for U.S. cement plants. Crucially, their interest margin is establishment entry/exit and dynamic investments in output capacity. However, they do not allow for input substitution. I show in Table 2.12 that a version of my model without input substitution yields an average elasticity between emission reduction and output more than half as large as in the full model, and closer to Fowlie et al. (2016).[18] Comparing this to the more flexible economy in which plants can substitute at both margins in Figure 2.11 sheds light on the critical role that input substitution plays in mitigating the loss of output for any emission reduction target.

Second, I show that heterogeneity in fuel productivity and inter-temporal switching between fuel sets also serve as a cautionary tale for larger-scale climate models in which an aggregate

[18] Note that a gap remains and the production frontier is still concave without input substitution at the plant level. This is for two reasons. First, even without input substitution, plants are differently affected by the tax based on their fuel sets, which affects aggregate input substitution due to the reallocation of output across plants. Second, the difference between the two elasticities is also attributed to the entry/exit margin in Fowlie et al. (2016), which decreases output and emissions through plants exiting in the aftermath of carbon policy.

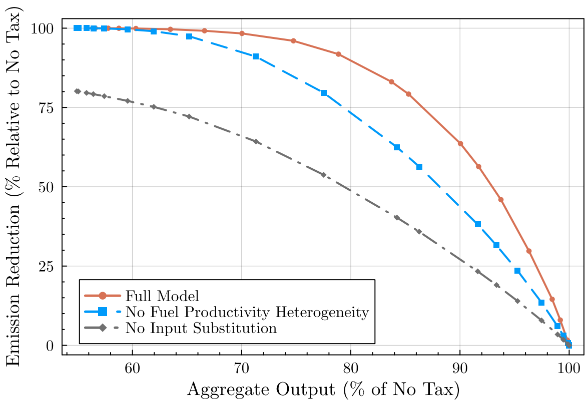

Table 2.11: Comparison of Trade-off Including No Input Substitution

	Average Elasticity $\frac{\%\Delta CO_{2e}}{\%\Delta Y}$
Full Model	5.46
No Fuel Productivity Heterogeneity	3.89
No Input Substitution	2.48
Fowlie et al. (2016) – U.S. Cement	1.04

Table 2.12: Comparison of Average Elasticity

Notes: The average elasticity of U.S. Cement plants is constructed by approximating Figure 2.A (aggregate output capacity) and 2.C (aggregate emissions) in Fowlie et al. (2016). They do various carbon policy exercises across different carbon prices. I specifically approximate their *Auctioning* policy, which is isomorphic to a carbon tax.

production function in different fossil fuels is typically assumed as part of a larger integrated assessment model (IAM). Such models study the relationship between climate change and economic growth. For example, Golosov et al. (2014) and Miftakhova and Renoir (2021) postulate an aggregate CES production for composite energy that combines different fuels. Such a production function may understate the extent of fuel substitution as a response to policy because it does not capture the reallocation mechanism induced by heterogeneity in fuel productivity and the inter-temporal substitution between fuel sets.

Technology Lock-in – Ineffectiveness of Carbon Tax at the Extensive Margin

While a carbon tax can cost-effectively reduce emissions by affecting fuel substitution at the intensive margin and reallocating output from high-emission to low-emission plants, its effect on the transition from coal to natural gas at the extensive margin is more limited. Indeed, any level of the carbon tax leads to a net decrease in the fraction of plants using coal or natural gas. The reduction in coal is relatively small, as it would take a carbon tax that raises the price of coal by 400% to incentivize a 10% decrease in the fraction of plants that use coal. Aside from coal being initially cheap relative to other fuels, this effect is primarily due to the option value that coal provides. Plants would rather reduce their coal consumption at the intensive margin but keep the option of using coal for the additional substitution margin it provides.

The lack of natural gas adoption indicates that the carbon tax is ineffective at incentivizing plants to overcome technology lock-in. This can be partially explained by a combination of the facts that the carbon tax also raises the price of natural gas, fixed costs of adoption are economically high – especially for plants away from the pipeline network, and plants who are not currently using natural gas would be on average 30% less productive at using natural gas

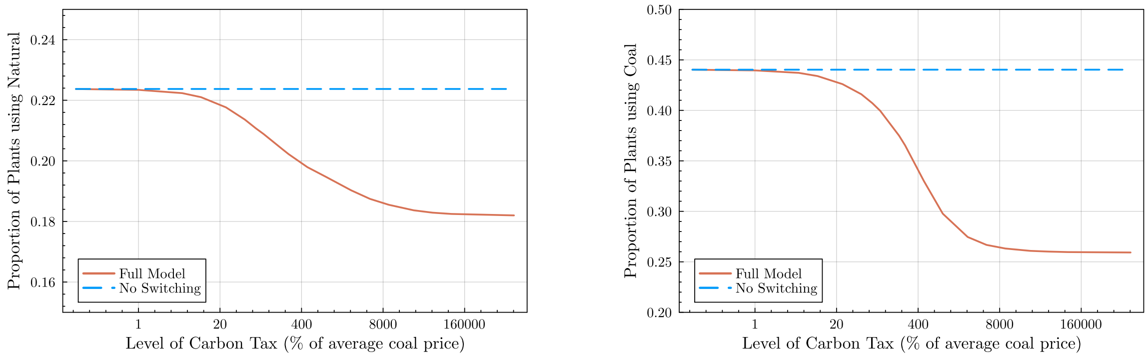

(a) Aggregate Share of Gas Users

(b) Aggregate Share of Coal Users

Figure 2.16: Natural Gas and Coal Take-up Across Levels of Carbon Tax

compared to those who already use it. In the next section, I explore the effectiveness of a subsidy to incentivize natural gas take-up.

2.8.3 Alleviating Technology Lock-in - Combining the Carbon Tax With Natural Gas Subsidy

To complement the carbon tax, I investigate how proceeds from the tax can be used to finance a subsidy to the fixed cost of natural gas to alleviate technology lock-in and increase natural gas take-up. Below, I show results for various permanent subsidies, ranging from 0% to 100% of the average fixed cost of natural gas. I do these experiments jointly with a carbon tax. To choose a representative social cost of carbon (SCC), I first set the social discount rate to 3% ($\beta = 0.97$) to match India's average real interest rate during the sampled period. Then, following the most recent estimates from the Inter-agency Working Group on the Social Cost of Carbon (IWG, 2021), I set the SCC to the 2020 estimates for a social discount rate of 3% at USD $51/tCO_{2e}$. This SCC corresponds to a mid-range estimate in the literature. In Figure 2.17, I show what happens to carbon tax revenues/externality damages and total subsidy paid out as the subsidy rate increases

First, up to 10% of the subsidy can be fully financed by a carbon tax in expectation. More importantly, carbon tax revenues monotonically increase as the subsidy rate increases. I show this more clearly in Figure 2.18, where I compare the evolution of the tax revenues with the fraction of plants that use natural gas and coal.

These results are important. As the subsidy rate increases, more plants add natural gas, but the fraction of plants that use coal remains almost the same. This isn't surprising because coal is still significantly cheaper than gas, and the salvage value of coal is much lower than gas. In

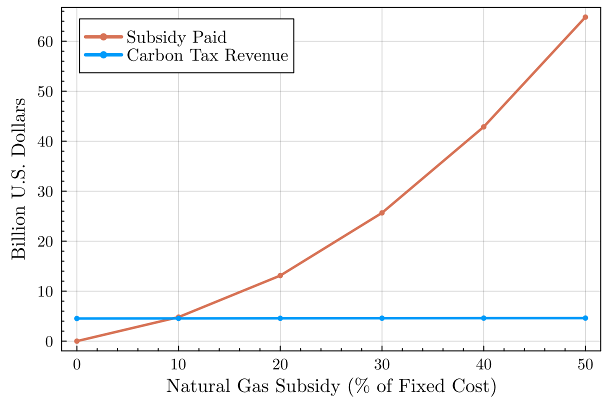

Figure 2.17: Carbon Tax Revenue and Subsidy Paid Out Along Subsidy Rate

Notes: This figure was calculated by simulating the expected total tax revenues and subsidies paid to plants for 40 years. Note that carbon tax revenues correspond to externality damages since carbon tax rates on each fuel are equal to marginal externality damages. As such, the evolution of the carbon tax revenues is indicative of emissions.

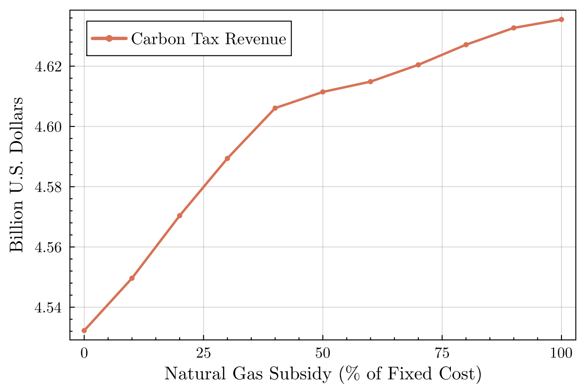

(a) Carbon Tax Revenues/Externality Damages

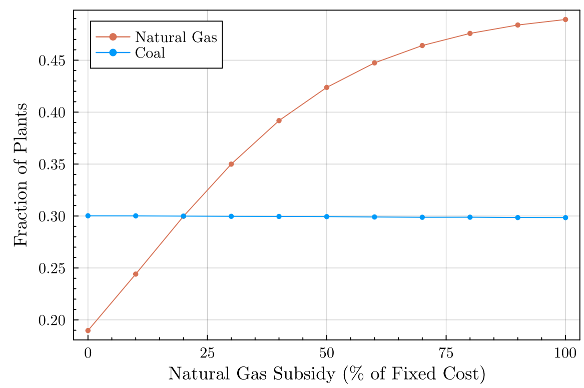

(b) Fraction of Plants who use Natural Gas/Coal

Figure 2.18: Comparison of Selected Outcomes Along Subsidy Rate

Notes: Carbon tax revenues correspond to externality damages since relative fuel tax rates are equal to relative fuel marginal externality damages. As such, the evolution of the carbon tax revenues is indicative of aggregate emissions.

this context, it makes more sense for plants to keep coal for the option value it provides. These results create two countervailing effects on carbon tax revenues/externality damages. On the one hand, plants that add natural gas substitute away from more polluting fuels such as oil, electricity, and coal. This substitution effect reduces tax revenues and externality damages. On the other hand, as plants add natural gas, they have more fuels available, increasing their option value and decreasing the price of energy. The net effect is a decrease in marginal production costs and an increase in output, which increases all input demand. Thus, even with substitution towards natural gas, the input demand for other fuels such as coal, oil, and electricity increases, increasing tax revenues and externality damages. In Figure 2.20, I do a full Shapley decomposition of the change in tax revenues as the subsidy rate increases between these two channels. Details of the Shapley decomposition can be found in Appendix D.2.

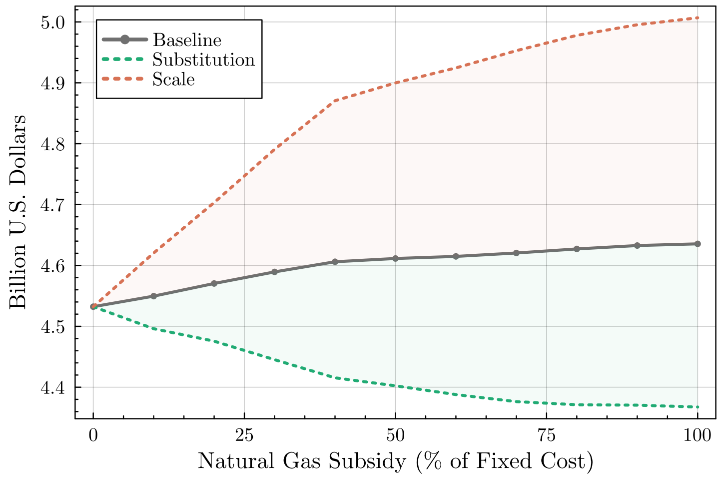

Figure 2.20: Shapley Decomposition of Changes in Tax Revenues

Unsurprisingly, the scale effect dominates. From a welfare perspective, it is unclear whether the subsidy is preferable to an economy with only a carbon tax. While both profits and consumer surplus increase, this comes at the cost of more externality damages and considerable investment subsidies, which could be allocated towards more profitable ventures. For this reason, I do a formal welfare analysis of this policy in the next section.

Welfare Analysis of a 10% Subsidy

I choose to narrow the focus on a 10% subsidy because it can be fully financed by the carbon tax, satisfying the government's budget constraint. With such a policy, per-period welfare is standard and features four components: consumer surplus, producer surplus, net government revenues, and externality damages (Fowlie et al., 2016):

$$w_t(\tau, s) = \underbrace{v_t(\tau, s)}_{\text{consumer surplus}} + \underbrace{\Pi(\tau, s)}_{\text{producer surplus}} + \underbrace{G(\tau, s)}_{\text{net gov. revenue}} - \underbrace{\sum_f \sum_i \gamma_f e_{fit}(\tau, s)}_{\text{externality damages}}$$

Where consumer surplus is decreasing in the aggregate output price index P_t. This is due to quasi-linear aggregate utility: $v_t(\tau, s) = \frac{\theta}{1-\theta} P_t(\tau, s)^{-\frac{\theta}{1-\theta}}$. As such, we can think of the remaining three parts of this welfare function as shifting the aggregate income of the consumers if it owns all plants and gets aggregate profits net of fixed costs, government revenues as lump-sum transfers, and suffers externality damages from pollution in dollars from the social cost of carbon. To include a subsidy towards natural gas adoption, I assume that the subsidy is financed by government revenue from the carbon tax and that every plant faces the same permanent subsidy amount of s. In this context, producer surplus is the sum of total profits net of subsidized fixed costs, and net government revenue is total tax revenues minus subsidy paid out.

$$\Pi(\tau, s) = \sum_{i=1}^{N} \Big(\underbrace{\pi_{it}(\tau, s)}_{\text{variable profits}} - \sum_{\mathcal{F}' \subseteq \mathbb{F}} \underbrace{[\mathcal{K}(\mathcal{F}' \mid \mathcal{F}_{it}) - s\mathbf{I}(gas \in \mathcal{F}' \setminus \mathcal{F}_{it})]\mathbf{I}(\mathcal{F}_{it+1} = \mathcal{F}' \mid \tau, s)}_{\text{subsidized fixed costs}} \Big)$$

$$G(\tau, s) = \sum_{i=1}^{N} \Big(\underbrace{\sum_{f} \tau_f e_{fit}(\tau, s)}_{\text{tax revenue}} - \sum_{\mathcal{F}' \subseteq \mathbb{F}} \underbrace{s\mathbf{I}(gas \in \mathcal{F}_{it+1} \setminus \mathcal{F}_{it})\mathbf{I}(\mathcal{F}_{it+1} = \mathcal{F}' \mid \tau, s)}_{\text{subsidy}} \Big)$$

Note that externality damages cancel out with tax revenue, and the subsidy cancels out because it is a transfer from $G(\tau, s)$ to $\Pi(\tau, s)$. As a result, period welfare is effectively equal to consumer surplus plus variable profits minus total fixed costs:

$$w_t(\tau, s) = \underbrace{v_t(\tau, s)}_{\text{consumer surplus}} + \underbrace{\sum_{i=1}^{N} \pi_{it}(\tau, s)}_{\text{variable profits}} - \underbrace{\sum_{i}^{N} \Big(\sum_{\mathcal{F}' \subseteq \mathbb{F}} \mathcal{K}(\mathcal{F}' \mid \mathcal{F}_{it})\mathbf{I}(\mathcal{F}_{it+1} = \mathcal{F}' \mid \tau, s) \Big)}_{\text{total fixed costs}} \tag{2.23}$$

Total welfare is then defined as the net present value of expected period welfare. I approximate total welfare by averaging multiple Monte-Carlo simulations of the economy (indexed by k) over a 40-year horizon. Lastly, the subsidy rate of s was chosen to keep the expected net government revenues weakly positive.

$$\begin{aligned} \mathcal{W}(\tau, s) &= \mathbb{E}_0 \Big(\sum_{t=0}^{\infty} \beta^t w_t(\tau, s) \Big) \qquad \mathbb{E}_0(G(\tau, s)) \geq 0 \\ &\approx \frac{1}{K} \sum_{k} \sum_{t=0}^{40} \beta^t \omega_{tk}(\tau, s) \end{aligned}$$

Below are the welfare results. In net, there is a small but positive welfare effect from the subsidy relative to a regime with only a carbon tax, which means using carbon tax revenues to subsidize the adoption of natural gas is slightly better than rebating it as a lump sum transfer to consumers. This welfare effect is explained by two countervailing effects. On one hand, variable profits and consumer surplus increased by 19 and 13 million dollars, respectively. This is because more plants add natural gas, but the fraction of plants using coal remains constant. This leads to a decrease in the average price of energy, a decrease in average marginal costs, and a decrease in output prices. Thus, more steel is produced at a lower cost, which benefits producers, and sold at a lower price, which benefits consumers. On the other hand, there is an increase in externality damages and an increase in total fixed costs paid in the economy by 10 and 31 million dollars, respectively. While externality damages cancel out with tax revenue, total fixed costs do not.

	Carbon Tax	**Carbon Tax + 10% Subsidy**	**Difference**
	Billion U.S Dollars	*Billion U.S Dollars*	*Million U.S. Dollars*
Total Welfare	63.415	63.417	1.18
Variable Profit	22.58	22.60	19.18
Consumer Surplus	22.21	22.22	13.62
Total Fixed Costs (*Paid by plants + subsidy*)	-18.633	-18.601	31.61
Externality Damages/Tax Revenue	2.64	2.65	9.78

Table 2.13: Decomposition of Welfare Effects – Carbon Tax With and Without Subsidy

Notes: All welfare components are reported by their net present value (NPV) over a horizon of 40 years from the last year of observation in the data (2016) with a social discount rate of 3%. Also, externality damages and tax revenue cancel out in the welfare function. The subsidy also cancels out because it is simply a transfer from net government revenue towards producer surplus. As a result, variable profits, consumer surplus, and total fixed costs are the remaining components in the welfare function such that $Welfare = ConsumerSurplus + VariableProfit - TotalFixedCost$

To understand how small the welfare effects are, I compare in Table 2.14 how much of the total fixed costs are financed by the subsidy. While 2.7 billion dollars go towards adopting natural gas, the fraction of plants that use natural gas only goes up by 20% from 0.18 to 0.24, while variable profits and consumer surplus jointly increase by 31 million dollars. Hence, private gains from the subsidy are only 1.1% of the policy's cost.

	Carbon Tax	**Carbon Tax + 10% Subsidy**	**Difference**
Fraction of Gas Users	0.19	0.24	0.05
Total Subsidy paid *(Billion U.S. Dollars)*	0.0	2.793	2.793

Table 2.14: Total Subsidy paid

Notes: This table reports the long-run fraction of plants that use natural gas after the policy and the net present value of the expected total subsidy paid to plants.

There are a few reasons explaining this small effect. First, by virtue of being a universal subsidy, the government effectively finances the adoption of natural gas for inframarginal plants that would have still adopted natural gas in the absence of the subsidy. This can be seen from the increase in total fixed costs paid in the economy by 31.61 million dollars in Table 2.13 after the policy, which is considerably lower than the total amount paid by the subsidy (2.7 billion). In practice, I find that over 90% of natural gas adopters are inframarginal. Second, plants at the margin that are incentivized to adopt natural gas in the aftermath of the policy are, on average, 30% less productive at using natural gas than plants that already use natural gas. See Figure 2.8. At the same time, natural gas is, on average, less productive per dollar invested into it than any other fuel. See Figure 2.15b.

This small welfare effect raises the question of whether the government could find more profitable avenues to invest proceeds from the carbon tax to alleviate technology lock-in. For example, it could invest in energy efficiency training programs to increase energy and fuel productivity or carbon capture technologies that reduce emissions ex-post. While outside the scope of this paper, this is an interesting avenue for future research.

2.9 Conclusion

In conclusion, I develop a rich dynamic production model to study fuel substitution from manufacturing establishments. It includes switching between fuel sets at a cost and heterogeneity in fuel productivity. By combining various methods from the production function estimation and the dynamic discrete choice literature, I show how this model can be estimated with a panel of plant-level data that features output and input prices/quantities. I then apply this model to the Indian Steel industry, which is high in energy and emission intensity due to the prevalence of coal usage. I then perform various counterfactual policy experiments to reduce emissions at the lowest cost possible, including a carbon tax and a carbon tax with a subsidy towards adopting cleaner fuels.

Moreover, I show that novel features of this model have important quantitative implications for the scope of these policies. Indeed, carbon taxation is much more targeted towards high-emission plants than previously thought due to multiple layers of heterogeneity. As a result, high-emission plants become relatively less competitive, reallocating output towards low-emission plants. This considerably reduces the overall economic cost of reducing emissions. However, more than a carbon tax is needed to increase adoption of cleaner fuels such as natural gas. For this reason, I show how proceeds from the carbon tax can be used to subsidize the fixed cost of natural gas adoption. There is a small but positive welfare effect, unexpectedly through a larger private surplus (producer and consumer) at the expense of higher emissions. This is due to the option value that an additional fuel provides, which lowers production costs. However, the welfare effects of the subsidy are minor compared to its cost. Overall, these results highlight the importance of producer heterogeneity and inter-temporal decisions when quantifying the impact of carbon policy.

Bibliography

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous, "The Environment and Directed Technological Change," *American Economic Review*, 2012, *102* (1), 131–166.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer, "Identification Properties of Recent Production Function Estimators," *Econometrica*, 2015, *83* (6), 2411–2451.

Ahmadi, Younes and Akio Yamazaki, "The Effectiveness of Revenue-Neutral Carbon Taxes," 2020.

Allcott, Hunt and Michael Greenstone, "Is There an Energy Efficiency Gap?," *Journal of Economic Perspectives*, 2012, *26* (1), 3–28.

Alpino, Matteo, Luca Citino, and Annalisa Frigo, "The effects of the 2021 energy crisis on medium-sized and large industrial firms: evidence from Italy," Questioni di Economia e Finanza (Occasional Papers) 776, Bank of Italy, Economic Research and International Relations Area June 2023.

Andersson, Julius J., "Carbon Taxes and CO2 Emissions: Sweden as a Case Study," *American Economic Journal: Economic Policy*, 2019, *11* (4), 1–30.

Arcidiacono, Peter and John Bailey Jones, "Finite mixture distributions, sequential likelihood and the EM algorithm," *Econometrica*, 2003, *71* (3), 933–946.

— **and Robert Miller**, "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

Atkinson, Scott E and Robert Halvorsen, "Interfuel Substitution in Steam Electric Power Generation," *Journal of Political Economy*, 1976, *84* (5), 959–978.

Atkinson, Scott E. and Rong Luo, "Estimation of Production Technologies with Output and Environmental Constraints," *International Economic Review*, 2023, *65*.

Blundell, Richard and Stephen Bond, "Initial conditions and moment restrictions in dynamic panel data models," *Journal of Econometrics*, 1998, *87*, 115–143.

_ **and** _ , "GMM Estimation with persistent panel data: an application to production functions," *Econometric Reviews*, 2000, *19* (3), 321–340.

_ **and** _ , "Initial conditions and Blundell-Bond estimators," *Journal of Econometrics*, 2023, *234*, 101–110.

Bollard, Albert, Peter J. Klenow, and Gunjan Sharma, "Indias mysterious manufacturing miracle," *Review of Economic Dynamics*, 2013, *16*, 59–85.

Bonhomme, Stéphane and Elena Manresa, "Grouped Patterns of Heterogeneity in Panel Data," *Econometrica*, 2015, *83* (3), 1147–1184.

Broda, Christian and David E Weinstein, "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, 2006, *121* (2), 541–585.

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer, "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?," *Journal of Policial Economy*, 2000, *108* (6).

Cho, Won G., Kiseok Nam, and José A. Pagan, "Economic growth and interfactor/interfuel substitution in Korea," *Energy Economics*, 2004, *26*, 31–50.

Christensen, Peter, Paul Francisco, and Erica Myers, "Can Incentive-Based Pay Increase the Marginal Value of Public Spending on Energy Efficiency ? Experimental Evidence from the Weatherization Assistance Program," 2022.

Collard-Wexler, Allan and Jan De Loecker, "Reallocation and Technology: Evidence from the US Steel Industry," *American Economic Review*, 2015, *105.*

Culler, Joseph A. and Eric T. Mansur, "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution," *American Economic Journal: Economic Policy*, 2017, *9* (3).

Demirer, Mert, "Production Function Estimation with Factor-Augmenting Technology: An Application to Markups," *Working Paper*, 2020.

Dhar, Subash, Minal Pathak, and Priyadarshi R. Shukla, "Transformation of India's steel and cement industry in a sustainable 1.5°C world," *Energy Policy*, 2020, *137.*

Doraszelski, Ulrich and Jordi Jaumandreu, "R and D and productivity: Estimating endogenous productivity," *Review of Economic Studies*, 2013, *80* (4), 1338–1383.

_ **and** _ , "Measuring the Bias of Technological Change," *Journal of Political Economy*, 2018, *126.*

Ethier, By Wilfred J, "National and International Returns to Scale in the Modern Theory of International Trade," *The American Economic Review*, 1982, *72* (3), 389–405.

Farmer, Leland E. and Alexis Akira Toda, "Discretizing nonlinear, non-Gaussian Markov processes with exact conditional moments," *Quantitative Economics*, 2017, *8.*

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, 2015, *105* (10), 3150–3182. available for download at `www.ggdc.net/pwt`.
Environmental Protection Agency - Center for Corporate Climate Leadership

Environmental Protection Agency - Center for Corporate Climate Leadership, "Emission Factors for Greenhouse Gas Inventories," 2022. data retrieved from the GHG Emission Factors Hub `https://www.epa.gov/system/files/documents/2022-04/ghg_emission_factors_hub.pdf`.
Indian Ministry of Steel

Indian Ministry of Steel, "Energy and Environmet Management in Iron & Steel Sector," 2023. Retrieved from `https://steel.gov.in/en/technicalwing/energy-and-environment-management-iron-steel-sector`.
Interagency Working Group on Social Cost of Greenhouse Gases

Interagency Working Group on Social Cost of Greenhouse Gases, United States Government, "Technical Support Document: Social Cost of Carbon, Methane, and Nitrous Oxide; Interim Estimates under Executive Order 13990," Technical Report February 2021.
United Nations — Climate Action

United Nations — Climate Action, "Causes and Effects of Climate Change," 2023. Retrieved from `https://www.un.org/en/climatechange/science/causes-effects-climate-change`.

Fontagné, Lionel, Philippe Martin, and Gianluca Orefice, "The Many Channels of Firm's Adjustment to Energy Shocks: Evidence from France," Working Paper 929, Banque de France Eurosystème 2023.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan, "Market-based emissions regulation and industry dynamics," *Journal of Political Economy*, 2016, *124* (1), 249–302.

Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker, "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes," *American Economic Journal: Applied Economics*, 2020, *12* (2), 303–342.

Gandhi, Amit, Salvador Navarro, and David A. Rivers, "On the identification of gross output production functions," *Journal of Political Economy*, 2020, *128* (8), 2973–3016.

Gittens, Jarone, "The Impact of Carbon Taxes on Market Structure: Evidence from British Columbia's Retail Gasoline Industry," Technical Report, Department of Economics, Queen's University 2020.

Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova, "Imported intermediate inputs and domestic product growth: Evidence from India," *Quarterly Journal of Economics*, 2010, *125* (4), 1727–1767.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinsky, "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 2014, *82* (1), 41–88.

Gosnell, Greer K., John List, and Robert D. Metcalfe, "The Impact of Management Practices on Employee Productivity: A Field Experiment with Airline Captains," *Journal of Political Economy*, 2020, *128* (4).

Goulder, Lawrence H. and Ian W. H. Parry, "Instrument Choice in Environmental Policy," *Review of Environmental Economics and Policy*, 2008, *2* (2).

Grieco, Paul L.E., Shengyu Li, and Hongsong Zhang, "Production function estimation with unobserved input price dispersion," *International Economic Review*, 2016, *57* (2), 665–690.

Gupta, Vaibhav, Tirtha Biswas, Deepa Janakiraman, and Karthik Ganesan, "Methodology Note - Greenhouse Gas Emissions of India - Subnational Estimates Industry Sector 2005 to 2015 series," Technical Report September 2019.

Harrison, Ann, Ben Hyman, Leslie Martin, and Shanthi Nataraj, "When do Firms Go Green? Comparing Price Incentives with Command and Control Regulations in India," 2016.

Hassler, John, Conny Olovsson, and Michael Reiter, "Integrated Assessment in a Multi-region World with Multiple Energy Sources and Endogenous Technical Change," 2019.

__ **, Per Krusell, Conny Olovsson, and Michael Reiter**, "On the effectiveness of climate policies," 2020, pp. 1–43.

Hawkins-Pierot, Jonathan T. and Katherine R. H. Wagner, "Technology Lock-In and Optimal Carbon Pricing," 2022.

Helpman, Elhanan and Oleg Itskhoki, "Labour Market Rigidities, Trade and Unemployment," *Review of Economic Studies*, 2010, *77* (3), 1100–1137.

Hyland, Marie and Stefanie Haller, "Firm-level Estimates of Fuel Substitution," *The Energy Journal*, 2018, *39* (6), 71–98.

Joskow, Paul L. and Frederic S. Mishkin, "Electric Utility Fuel Choice Behavior in the United States," *International Economic Review*, 1977, *18* (3), 719–736.

Kasahara, Hiroyuki and Joel Rodrigue, "Does the use of imported intermediates increase productivity? Plant-level evidence," *Journal of Development Economics*, 2008, *87* (1), 106–118.

Klette, Tor Jakob and Zvi Griliches, "The inconsistency of common scale estimators when output prices are unobserved and endogenous," *Journal of Applied Econometrics*, 1996, *11* (4), 343–361.

León-Ledesma, Miguel A., Peter McAdam, and Alpo Willman, "Identifying the elasticity of substitution with biased technical change," *American Economic Review*, 2010, *100* (4), 1330–1357.

Levinsohn, James and Amil Petrin, "Estimating production functions using inputs to control for unobservables," *Review of Economic Studies*, 2003, *70* (2), 317–341.

Levinson, Arik, "A Direct Estimate of the Technique Effect: Changes in the Pollution Intensity of US Manufacturing, 1990–2008," *Journal of the Association of Environmental and Resource Economists*, 2015, *2* (1).

Lyubich, Eva, Joseph Shapiro, and Reed Walker, "Regulating Mismeasured Pollution: Implications of Firm Heterogeneity for Environmental Policy," *AEA Papers and Proceedings*, 2018, *108*, 136–142.

Ma, Hengyun, Les Oxley, John Gibson, and Bonggeun Kim, "China's energy economy: Technical change, factor demand and interfactor/interfuel substitution," *Energy Economics*, 2008, *30*, 2167–2183.

Miftakhova, Alena and Clément Renoir, "Economic Growth and Equity in Anticipation of Climate Policy," *SSRN Electronic Journal*, 2021, (June).

Miller, Jack Robert, "The Direct Reduction of Iron Ore," *Scientific American*, 1976, *235* (1), 68–81.

Nordhaus, William, *A Question of Balance: Weighing the Options of Global Warming Policies*, Vol. 22, New Haven London: Yale University Press, 2008.

Oberfield, Ezra and Devesh Raval, "Micro Data and Macro Technology," *Econometrica*, 2021, *89* (2), 703–732.

Olley, G. Steven and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 1996, *64* (6), 1263–1297.

Pindyck, Robert S., "Interfuel Substitution and the Industrial Demand for Energy," *The Review of Economics and Statistics*, 1979, *61* (2), 169–179.

Ramanarayanan, Ananth, "Imported inputs and the gains from trade," *Journal of International Economics*, 2020, *122*, 103260.

Rehfeldt, Matthias, Tobias Fleiter, Andrea Herbst, and Stefan Eidelloth, "Fuel switching as an option for medium-term emission reduction - A model-based analysis of reactions to price signals and regulatory action in German industry," *Energy Policy*, 2020, *147* (September), 111889.

Romer, Paul M, "Endogenous Technological Change," *Journal of Political Economy*, 1990, *98* (5).

Scott, Jonathan B., "Positive Spillovers from Infrastructure Investment: How Pipeline Expansions Encourage Fuel Switching," *The Review of Economics and Statistics*, 2021, pp. 1–43.

Shapiro, Joseph S. and Reed Walker, "Why is pollution from us manufacturing declining? the roles of environmental regulation, productivity, and trade," *American Economic Review*, 2018, *108* (12), 3814–3854.

Wang, Xiaolei and Boqiang Lin, "Factor and fuel substitution in China's iron & steel industry: Evidence and policy implications," *Journal and Cleaner Production*, 2017, *141*, 751–759.

Worrell, Ernst, Lenny Bernstein, Joyashree Roy, Lynn Price, and Jochen Harnisch, "Industrial energy efficiency and climate change mitigation," *Energy Efficiency*, 2009, *2* (2), 109–123.

Zhang, Hongsong, "Non-neutral technology, firm heterogeneity, and labor demand," *Journal of Development Economics*, 2019, *140* (May), 145–168.

Chapter 3

Asymmetric Carbon Pricing, Fuel Substitution and Carbon Leakage

3.1 Introduction

Many countries are implementing regulations to combat pollution and its associated environmental challenges. As a result, there is an increasing interest in assessing the effectiveness of these policies in reducing pollution. The combustion of fossil fuels, a major contributor to pollution, creates a negative externality through the emission of greenhouse gases that contribute to global warming. Under Pigouvian theory, the optimal policy for addressing this externality is a carbon tax equal to the marginal social damages of GHG emissions. Although many countries and regions have implemented carbon taxes, the limited jurisdictional scope of regulation is such that even the most ambitious pollution reduction programs, such as the Kyoto Protocol and the Paris Agreement, are voluntary, and their effectiveness depends on the goodwill of governments. In this context, there is a risk of "carbon leakage", where emissions shift to unregulated regions as a result of asymmetric regulation.

This risk of carbon leakage is particularly relevant in the context of manufacturing activity because firms are known to compete across regions (Smith and Ocampo, 2020). Moreover, manufacturing activity contributes to 37% of global greenhouse gas emission (Worrell, Bernstein, Roy, Price and Harnisch, 2009) and is one of the main targets of carbon taxes.

Concerns about carbon leakage are accentuated in the case of sub-national and even sub-union regulation because production is more likely to shift across borders within a country/union due to limited trade barriers. Moreover, while carbon leakage is by itself an interesting phenomenon, its presence biases empirical studies that estimate the direct effect of carbon taxes. Indeed, the competitive nature of firms across regions makes it more difficult to exploit the jurisdictional boundaries of policies as natural experiments to form counterfactuals because it implies that the policy also treats unregulated firms. To use standard policy evaluation tools such as Difference-in-Difference to recover the effect of a policy, one must then make strong independence assumptions that rule out carbon leakage.

In this paper, I study carbon leakage in the context of the British Columbia (B.C.) and Quebec carbon taxes implemented in 2007 and 2008, respectively. I build a model of monopolistic competition with heterogeneous firms and multiple regions similar to Shapiro and Walker (2018) and Aichele and Felbermayr (2015). In the model, a carbon tax increases regulated firms' marginal costs, making them less competitive than unregulated firms. The increase in marginal costs depends on firms' capacity to substitute cleaner fuels for dirty fuels, and the extent of carbon leakage depends on how consumers are willing to substitute across firms. The model allows me to study both the direct effects of the carbon tax in regulated regions and the leakage effect in unregulated regions. Most importantly, I provide restrictions to quantify this model—and, by extension, carbon leakage–using publicly available emissions data/remote sensing data only.

Using publicly available Canadian data on a wide range of pollutants emitted in the air by manufacturing establishments, I estimate the model's parameters to quantify the impact of the British Columbia (B.C.) and Quebec carbon taxes on greenhouse gas (GHG) emissions across Canada prior to the uniform Canadian carbon tax implemented in 2018. Both carbon taxes were implemented in 2007 and 2008, respectively. I find strong evidence of carbon leakage to other Canadian provinces. Indeed, I find that the carbon taxes caused a decrease in aggregate emissions from Quebec and B.C. firms by 150 megatons of carbon dioxide equivalent (CO_{2e}) from 2007 to 2015. However, these carbon taxes caused an increase in emissions from firms in unregulated provinces by 67 megatons of CO_{2e}, mitigating 45% of emissions reduction efforts as output reallocated towards unregulated firms.

I find very little evidence of fuel substitution in response to the tax. Most firms use a dominant fuel and do not find substitution profitable. Moreover, the dominant fuel tends to be natural gas, whereas alternative fuels such as oil and coal are even more emission-intensive. As a result, B.C. firms pass on the cost increase to consumers and reduce output. This cost pass-through makes firms in unregulated provinces relatively more competitive, increasing their market share.[1]

[1]However, I do not observe electricity consumption by firms, which could mitigate this leakage effect if regulated firms can substitute electricity for fossil fuels, especially in provinces such as Quebec, where electricity is "clean" as it comes mostly from hydro plants.

The counterfactual introduction of a uniform carbon tax across all provinces dramatically reduces aggregate emissions[2]. Extending the carbon tax to all provinces reduces aggregate emissions by 21% relative to a 2.15% reduction in aggregate emissions from the B.C. and Quebec carbon taxes.

The model features monopolistic competition in the spirit of Melitz (2003) and multiple regions. In a nontrivial extension of Copeland and Taylor (2004), I allow multiple fossil fuels to be used in production and specify pollution as a byproduct of fossil fuel combustion. Fossil fuel combustion generates energy for production but also releases greenhouse gases (GHG) into the atmosphere. This production function directly maps to emissions and carbon tax data. Indeed, I calculate quantities of fossil fuels from emissions data based on underlying chemical reactions. Carbon taxes are then levied with a different per-unit tax rate on different fuels, reflecting the emission intensity of each fuel, which affects firms' cost of using different fuels[3].

This extension allows me to identify and estimate the model's parameters and provide novel empirical evidence on carbon leakage using only publicly available pollution release data, fuel price data, and industry aggregates, all commonly available. The key to my method is that, under a standard set of assumptions about firms' production process and profit maximization, pollution data contains information about which fuels firms use, how much of each fuel they use, and how productive they are relative to other firms in the same industry.

My method applies to contexts with abundant remote sensing data but where plant-level data is difficult to access or hardly even exists. It is also one of the first papers to empirically study carbon leakage at a granular level in a literature that has a longstanding theoretical foundation (Garella and Trentinaglia, 2019, Hoel, 1991, Holland, 2012) and computational general equilibrium (CGE) applications (Böhringer, Carbone and Rutherford, 2016, Felder and Rutherford, 1993). Moreover, since carbon tax rates are heterogeneous by fuel types, I allow firms to substitute across different fuels. Similar to the approaches of Atkinson and Luo (2023), Ganapati, Shapiro and Walker (2020), and Carlson, Burtraw, Cropper and Palmer (2000), fuel substitution provides an endogenous margin of pollution intensity adjustment. Interestingly, I find very little evidence of fuel substitution induced by the B.C. and Quebec carbon taxes.

I exploit carbon taxes and plausibly exogenous variation in world fuel prices to identify the model's parameters. I assume that the production technology features constant elasticity of substitution (CES), and I separate identification into two parts. The first part comes from firms' cost minimization problem of choosing relative fuel shares to form a composite fuel index. To

[2]Such as tax was implemented in 2018 by the Federal Government.

[3]This is how many carbon taxes are implemented in practices, including the B.C. carbon tax.

this end, I show closed-form identification of reduced-form parameters that directly map into structural technology parameters: fuel-specific efficiency parameters that vary across industries and the elasticity of substitution across fuels. The second part comes from firms' profit maximization problem, in which they choose output quantity subject to a given level of the composite fuel index. I now observe an estimate of firms' fuel composite input, which contains information about their productivity. Under parametric assumptions about the productivity distribution, I use maximum likelihood to estimate the elasticity of substitution across firms, directly leveraging variation in the B.C. and Quebec carbon tax.

The data I use is the National Pollutant Release Inventory (NPRI) combined with industry-level aggregates. The former contains all pollutants released by manufacturing plants in Canada. I specifically use pollutants released in the air that are known to come from the combustion of specific fossil fuels. I invert each fuel's chemical reaction to recover an estimate of fuel quantities from a range of pollutants. This procedure is key because the model relies on fuel quantities to identify the effect of carbon taxes.

Section 3.2 presents the intuition underlying carbon leakage. Section 3.3 presents the model in all its details. Section 3.4 presents the data. Section 3.5 presents identification and details on estimation. Section 3.6 presents the results on estimation. Section 3.7 presents the counterfactual of interests and decomposes the effect of carbon taxes.

3.2 Carbon Leakage Mechanisms

In this section, I explain intuitively the different mechanisms through which an asymmetric carbon tax can lead to carbon leakage. Theoretically, there are two mechanisms at the intensive margins (within-firm adjustments) and one at the extensive margins (through entry and exit). First, the carbon tax increases the marginal cost of regulated firms, which leads to an increase in their output price and a reallocation of production across regions due to demand shifting from regulated to unregulated firms. Second, the increase in marginal costs can increase the productivity threshold required for a firm to be profitable, forcing some regulated firms to exit the market and vice versa in unregulated regions with firms entering the market. Third, the carbon tax also increases the relative price of most polluting fuels, and regulated firms will increase their demand for cleaner fuels. In general equilibrium, such variation in fuel demand will increase the relative gross price of cleaner fuels. Unregulated firms now face lower relative prices of polluting fuels, inducing substitution towards these polluting fuels.

I highlight the last two channels in red In Figure 3.1, indicating that these channels are not a priori relevant for Canadian firms in response to the B.C. and Quebec carbon taxes. Canadian

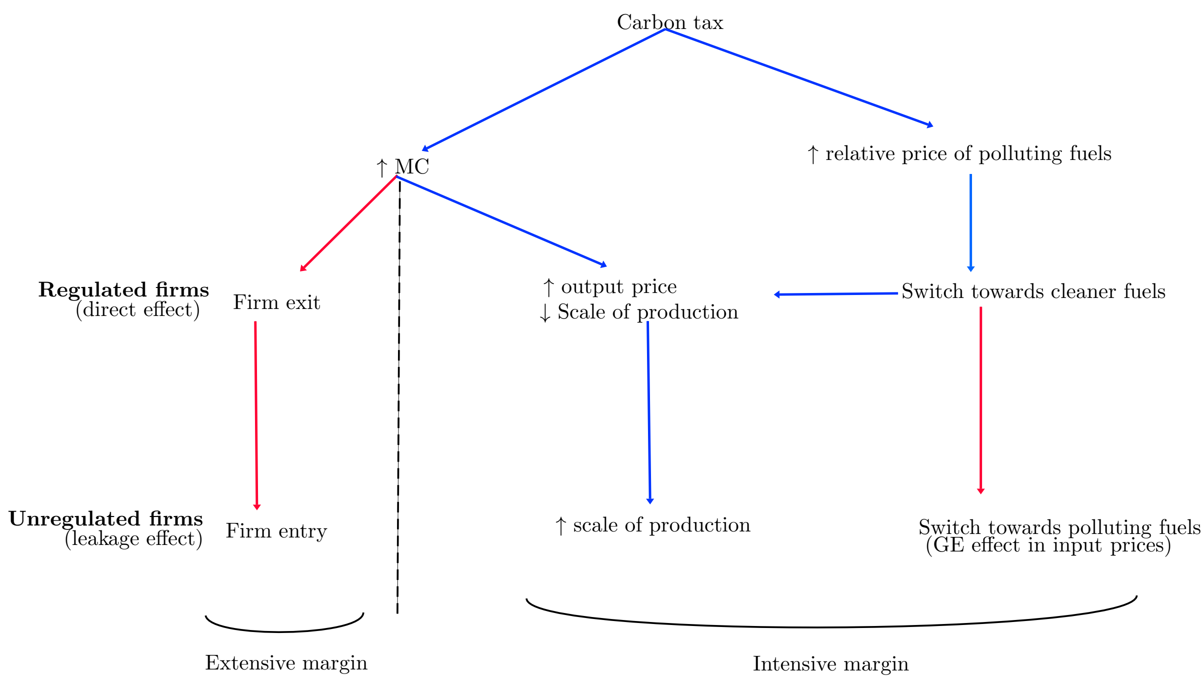

Figure 3.1: Channels for Carbon Leakage

firms are too small relative to the rest of the world to significantly impact aggregate fuel demand and induce change in world fuel prices. Moreover, I abstract from the entry/exit margin of adjustment as preliminary evidence suggests against it. Indeed, difference-in-difference (DiD) estimation on the number of firms operating in each province due to the B.C. and Quebec carbon taxes suggests an increase in the number of firms operating in both regulated provinces, which is at odds with theoretical predictions. A decomposition of this DiD estimator indicates a significant decrease in the number of firms operating in unregulated regions in the post-tax periods (see Appendix for details). For these reasons, I only look at the channel highlighted in blue, even though the model allows for all three mechanisms.

3.3 Model

3.3.1 Structure of the Economy

The economy is characterized by firms who engage in monopolistic competition across multiple regions and industries and share similarities with Shapiro and Walker (2018) and Melitz and Redding (2014). I augment this framework with firm-specific production functions that take different fuels as inputs and allow for inter-fuel substitution. I introduce an asymmetric carbon tax that only affects firms in specific regions to this economy. I first present the main framework for a single region and will give details on multiple regions when introducing the carbon tax and carbon leakage. A Cobb Douglas production function takes the output of J industries and aggregates them into a final good:

$$Y = \prod_{j=1}^{J} Y_j^{\beta_j} \text{ with } \sum_j \beta_j = 1$$

Where β_j is the share of industry j's production allocated to final consumption. The final good producer chooses how much of each industry's aggregate is needed to maximize some fixed amount of final consumption subject to a standard budget constraint: $\sum_j P_j Y_j = C$ where C denotes aggregate income. Throughout the paper, I assume that aggregate income is in dollar units and that firms take it as given. The solution to this problem is standard and yields the share of total consumption produced by each industry, Y_j:

$$Y_j = \frac{\beta_j}{P_j} C \tag{3.1}$$

Where P_j is the industry price index. Next, within each industry, firms sell differentiated goods, $Y_j(A)$, indexed by productivity A, which can be substituted at rate $\rho > 1$ to form the CES composite industry output Y_j. ρ will be a crucial parameter to study carbon leakage as it defines the degree of competitivity across firms and how easily consumers can switch between firms.

$$Y_j \equiv \left(\int_{\Omega_j} Y_j(A)^{(\rho-1)/\rho} dA \right)^{\rho/(\rho-1)}$$

The difference in integrating regions Ω_j represents the idea that different industries have different masses of operating firms indexed by N_j. From this, I can solve for the share of each firm's output $Y_j(A)$ as a fraction of the industry aggregate Y_j, which is easily found by choosing the cost-minimizing bundle of $\{Y_j(A)\}$ that produces a given amount of Y_j taking firm-level prices as given:

$$Y_j(A) = \left(\frac{P_j(A)}{P_j}\right)^{-\rho} Y_j \tag{3.2}$$

Where $P_j \equiv \left(\int_{\Omega_j} P_j(A)^{(1-\rho)} dA \right)^{1/(1-\rho)}$ is the industry-specific CES price index. Moving forward, I assume that firms compete in monopolistic competition.

3.3.2 Technology and Emissions

To the standard framework above, I add a CES technology that takes multiple fuels $\{q^\ell\}_{\ell=1}^L$ indexed by ℓ as inputs which form a fuel composite index F_j. This composite fuel represents the total quantity of energy service received from fuel combustion, which also emits greenhouse gases in the air. This technology is similar to the aggregate production function in Hassler, Olovsson and Reiter (2019), who study multiple energy sources in an integrated assessment model of climate change. Fuels have varying degrees of emission intensity, and firms can substitute between them at a rate $\sigma > 1$. Each industry j has baseline fuel-specific efficiency terms, $\lambda_{\ell j}$, where $\sum_\ell \lambda_{\ell j} = 1$. Each firm has a productivity level A drawn from a distribution $g_j(A)$. Productivity should be considered energy-augmenting rather than total factor productivity because it contains unobserved factors of production like labor and capital.

$$Y_j(A) = A \underbrace{\Big(\sum_\ell \lambda_{\ell j} (q_j^\ell)^{(\sigma-1)/\sigma} \Big)^{\sigma/(\sigma-1)}}_{F_j} \tag{3.3}$$

Let γ_ℓ be the coefficient that maps one unit of fuel ℓ to tons of carbon dioxide equivalent (CO_{2e}) equivalent, the measure for GHG emissions commonly used in the literature. Then, GHG emissions of firm i in industry j is the sum of all fuels used times their emission factor, which I decompose into the product of emission intensity and output quantity:

$$\begin{aligned} GHG_j(A) = \sum_\ell \gamma_\ell q_j^\ell(A) \equiv E_j(A) &= \frac{E_j(A)}{Y_j(A)} \times Y_j(A) \\ &= \underbrace{e_j(A)}_{\text{emission intensity (process factor)}} \times \underbrace{Y_j(A)}_{\text{output quantity (scale factor)}} \end{aligned}$$

It is useful to compare my technology to the canonical technological framework for pollution in the literature (Levinson and Taylor, 2008, Shapiro and Walker, 2018). I define pollution through a composite fuel index F_j, which is the CES aggregate of multiple fuels, whereas they define an implicit pollution function:

$$\begin{aligned} \text{My model: } & Y_j(A) = AF_j \\ \text{Canonical model: } & Y_j(A) = (Al_j)^{1-\alpha_j}(z_j)^{\alpha_j} \end{aligned}$$

In the canonical model, l_j is labor, and z_j is total pollution such that one unit of z_j always pollute the same amount. However, firms can have varying emission intensity through investments in pollution abatement technologies, making pollution less intensive relative to labor (lower α_j). By contrast, I define pollution abatement endogenously through fuel substitution. One unit of F_j can emit different levels of GHG emissions depending on the underlying bundle of fuels that compose it, and that will create variation in emission intensity. A specification that maps variation in energy intensity to fuel substitution is desirable in the context of greenhouse gas emissions, where most emissions reductions come from substituting cleaner fuels for dirty fuels rather than the end-of-pipe solutions. Moreover, I can exploit the mapping between the inputs (fossil fuels) and emissions to study emissions reduction efforts.

This technology also has another important implication. By introducing fuel-specific efficiency terms $\lambda_{\ell j}$, firms in different industries face different realized fuel prices in terms of energy service per dollar. I call these prices "effective prices", even though underlying observed fuel prices (in units of fuel quantity per dollar) may be the same. This heterogeneity in effective prices implies that a firm with higher relative efficiency in a specific fuel perceives it as cheaper than other fuels because it can use it better. This is motivated by the empirical fact that firms in different industries purchase very different relative fuel quantities even though they often face the same fuel prices.

$$\textbf{Observed}: \begin{bmatrix} p_\ell \\ p \end{bmatrix} = \begin{bmatrix} p_{\ell s} \\ \left(\sum_\ell p_\ell^{1-\sigma} \right)^{1/(1-\sigma)} \end{bmatrix} \longrightarrow \textbf{Effective:} \begin{bmatrix} \tilde{p}_{\ell j} \\ \tilde{p}_j \end{bmatrix} = \begin{bmatrix} \frac{p_\ell}{\lambda_{\ell j}} \\ \left(\sum_\ell \lambda_{\ell j}^\sigma p_\ell^{1-\sigma} \right)^{1/(1-\sigma)} \end{bmatrix}$$

Moreover, due to the constant returns assumption, individual effective fuel prices can be aggregated into an effective fuel price index $\tilde{p}_j$ that buys quantities of the composite fuel index F_{ij}, which is akin to labor as a single input in Melitz (2003)[4]. This makes the analysis of firms' scale decisions separate from the analysis of relative input choices. These features have important theoretical implications and will greatly simplify the model estimation.

Asymmetric carbon tax

[4]I assume constant returns to scale in the production of the fuel composite. However, this does not require constant returns in other inputs such as labor, capital, and intermediate materials. Here, other inputs get absorbed by productivity A because they are unobserved. If these inputs were observed, the constant returns fuel production function could be nested into another production function that takes the fuel composite along with labor, capital, and intermediate materials as inputs. This is the approach typically taken in the fuel substitution literature (Atkinson and Luo, 2023, Cho et al., 2004, Hyland and Haller, 2018, Ma et al., 2008, Pindyck, 1979, Wang and Lin, 2017).

I now introduce the main object of analysis to this framework: a carbon tax that affects a fraction $\tilde{N}_j^r$ of firms within each industry. Regulated firms are indexed by $s = r$ and face an additional per-unit tax rate $\{\tau_\ell\}_{\ell=1}^{L}$ which is added to gross fuel prices $p_{\ell r} = p_\ell + \tau_\ell \; \forall \ell$, while the remaining fraction $\tilde{N}_j^u$ unregulated firms are indexed by status $s = u$ and face same gross fuel price as before. Since this is a carbon tax, the fuel-specific tax rate is weakly increasing in fuel pollution intensity:

$$\gamma_\ell \geq \gamma_k \rightarrow \tau_\ell \geq \tau_k \; \forall k \neq \ell$$

The rationale behind this framework is that implementing a uniform tax rate on GHG emissions is typically achieved with a different tax rate across fuels due to varying fuel emission intensity γ_ℓ. This is exactly how B.C. introduced its carbon tax. For example, coal combustion emits, on average, twice as much CO_{2e} in the air as natural gas, and its tax rate is twice as high.

3.3.3 Firms' Optimal Decisions

Output Quantity and Price

By the constant returns to scale assumption, marginal costs are constant, and I can solve firms' problems in two parts which correspond to the scale factor and the process factor of GHG emissions. This is important because it allows me to quantify the model using pollution release data only. First, I solve the profit-maximizing amount of output quantity $Y_{js}(A)$ that a firm with productivity A in industry j and regulation regime s wants to produce from purchasing the composite fuel good $F_{js}(A)$, taking as given the composite effective fuel price index $\tilde{p}_{js}$. From this, I can know the equilibrium output of each firm $Y_{js}(A)$, the share of output allocated to fuels $F_{js}(A)$, and the output price $P_{js}(A)$, which will all be reflected in the scale factor of GHG emission. In the second part, I can solve the cost-minimizing bundle of fuels that will form $F_{js}(A)$. As such, I map the relative share of each fuel bought to the emission intensity of each firm, which forms the process factor of GHG emission.

I assume that firms compete over quantity and face an inverse demand derived from Equations 3.1 and 3.2. To avoid notational clutter, I will do the exposition for a single industry and remove the j subscript. Then, taking as given industry aggregates P and Y, a firm with productivity A in regulation status s solves:

$$\max_{Y_s(A)} \left\{ P_s(Y_s(A))Y_s(A) - \tilde{p}_s F_s(A) \right\}$$
$$\text{s.t. } P_s(Y_s(A)) = \left(\frac{Y_s(A)}{\beta C}\right)^{-1/\rho} P^{(\rho-1)/\rho}$$

Since $F_s(A) = \frac{Y_s}{A}$, each firm's marginal cost is the ratio of input price index to productivity: $c_s(A) = \frac{\tilde{p}_s}{A}$. This marginal cost leads to a standard monopolistic competition equilibrium pricing equation of a constant markup over marginal costs:

$$P(\tilde{p}_s, A) = \frac{\rho}{\rho - 1}\frac{\tilde{p}_s}{A} \tag{3.4}$$

Aggregate Productivity

This individual pricing equation can be aggregated into an industry price index as in Melitz (2003). To do so, I assume that productivity comes from a known distribution with density $g(A)$. In principle, this productivity distribution could be used to derive the productivity distribution of active firms in each region, $\mu_s(A)$, which would depend on a productivity threshold that defines which firms operate in the market. This threshold would depend on regulation status since the carbon tax raises the relative marginal cost of regulated firms, thus increasing the productivity required for the marginal regulated firm to operate and decreasing the productivity required for the marginal unregulated firm. However, in the Appendix, I show that entry/exit does not appear to be an empirically relevant margin of adjustment in the context of the B.C./Quebec carbon taxes, which is why I shall assume for the remainder of this exposition that the productivity threshold is the same across regulation status $\mu_s(A) = \mu(A)$ and I assume an interior solution for all firms such that the distribution of productivity for active firms is the same as the distribution of productivity for all firms, $\mu(A) = g(A)$. Then,

$$\tilde{A} = \Big(\int_0^\infty A^{\rho-1} g(A)dA \Big)^{1/(\rho-1)} \tag{3.5}$$

With a share $\tilde{N}^r = N^r/N$ of regulated firms and $\tilde{N}^u = N^u/N$ of unregulated firms, the equilibrium industry price index will be a CES aggregate of regulated and unregulated individual firm prices. Moreover, when $g_s(A) = g(A)$, this price index has the same structure as firm-level output prices, where aggregate marginal cost is a combination of regulated and unregulated input prices weighted by the respective mass of firms in each status over aggregate productivity:

$$
\begin{aligned}
P &= \Big(\sum_s \tilde{N}^s \int_0^\infty P_s(A)^{1-\rho} g(A) dA\Big)^{1/(1-\rho)} \\
&= \frac{\rho}{\rho-1}\Big(\tilde{p}_r^{1-\rho}\tilde{N}^r + \tilde{p}_u^{1-\rho}\tilde{N}^u\Big)^{1/(1-\rho)}\Big(\int_0^\infty A^{\rho-1} g(A) dA\Big)^{1/(1-\rho)} \\
&= \frac{\rho}{\rho-1}\frac{\Big(\tilde{p}_r^{1-\rho}\tilde{N}^r + \tilde{p}_u^{1-\rho}\tilde{N}^u\Big)^{1/(1-\rho)}}{\tilde{A}}
\end{aligned}
\tag{3.6}
$$

Note that absent of the carbon tax, this price index collapses to the price index in Melitz (2003) where labor is replaced by the composite fuel index: $P = \frac{\rho}{\rho-1}\frac{\tilde{p}}{\tilde{A}}$. Putting everything together, I can now find the quantity of the composite fuel purchased by a firm with productivity A and regulation status s:

$$
F_s(A) = \frac{\rho-1}{\rho}\frac{\tilde{p}_s^{-\rho}}{(\tilde{p}_r^{1-\rho}N^r + \tilde{p}_u^{1-\rho}N^u)}\beta C\Big(\frac{A}{\tilde{A}}\Big)^{\rho-1} \tag{3.7}
$$

The amount of composite fuel a firm wants decreases in the fuel price index and increases in productivity due to a competitive effect that gives productive firms a larger sales share due to their cost advantage. Since I assume that aggregate consumption is Cobb-Douglas across industries, there is no substitution across industries (net of the income effect), and only individual productivity relative to the industry average matters for firms' decisions.

relative fuel share and input price index

To choose the cost-minimizing share of each fuel that composes $F_s(A)$, firms face the technology defined in equation 3.3, and their relative fuel choice will only be a function of the interior parameters of the technology, namely interfuel substitutability σ and fuel efficiencies λ_ℓ. Productivity will not matter for relative fuel quantities because it augments the composite fuel index rather than specific fuels. Moreover, I assume that firms take input prices as given and cannot affect such prices with their decisions because they are too small relative to the population of firms that make up global fuel demand. This is motivated by the idea that the set of Canadian manufacturing firms is small relative to the set of all firms and consumers that buy fuels around the world. Additionally, supply-side shocks such as new technologies (e.g., fracking) and geopolitical events that shift fuel supply in specific regions often drive fuel prices. Therefore, taking input prices and F as given, a firm in regulation status s solves the following:

$$\min_{\{q_s^\ell\}_{\ell=1}^L} \Big\{ \sum_\ell p_{\ell s} q_s^\ell \Big\} \text{ s.t. } F = \Big(\sum_\ell \lambda_\ell (q_s^\ell)^{(\sigma-1)/\sigma} \Big)^{\sigma/(\sigma-1)}$$

The solution to this problem gives rise to the perceived input price index that varies across regulation status and fuel-specific perceived prices, whose implications were discussed earlier: $\tilde{p}_s = \frac{p_{\ell s}}{\lambda_\ell}$:

$$\tilde{p}_s = \Big(\sum_\ell \lambda_\ell^\sigma p_{\ell s}^{1-\sigma} \Big)^{1/(1-\sigma)}$$

This perceived price index is higher for regulated firms because $p_{\ell r} = p_\ell + \tau_\ell$ while $p_{\ell ur} = p_\ell$ and $\sigma > 1$. With this, I can define the share of fuel ℓ that makes up the composite fuel F, which is a simple function of relative perceived input prices. Indeed, as the perceived price of a fuel increases, the relative quantity of that fuel decreases at rate σ due to substitution towards other fuels.

$$q_s^\ell(F) = \Big(\frac{\tilde{p}_{\ell s}}{\tilde{p}_s} \Big)^{-\sigma} F \tag{3.8}$$

Emission intensity

Emission intensity is the amount of GHG emissions as a fraction of output that a firm produces and can be endogenously determined by the conditional input demand of each fuel times its emission factor $\gamma_{\ell;}$ over the firm's output quantity:

$$\begin{aligned} e_s(A) &= \frac{\sum_\ell \gamma_\ell q_s^\ell(F_s)}{Y_s(A)} \\ &= \frac{1}{A} \sum_\ell \gamma_\ell \bigg(\frac{\tilde{p}_{\ell s}}{\tilde{p}_s} \bigg)^{-\sigma} \end{aligned}$$

As in Shapiro and Walker (2018), emission intensity is locally decreasing in productivity. Moreover, while there is no capital and labor in the technology, if a firm is highly intensive in capital and/or labor relative to fuel, this model will capture this as higher productivity (less fuel required to produce a unit of output), hence a lower emission intensity, which is precisely what would happen if capital and labor were in the model[5].

When a firm gets regulated, it faces higher prices to purchase fuels that pollute more relative to other fuels. Due to possible substitution, it will change its optimal combination of fuels that make up F, and the distribution will shift towards fuels with lower emission factors, and emission intensity will decrease. For a set of fixed input prices, no carbon leakage is induced by emission intensity variation because unregulated firms still face the same gross price as before and will choose the same optimal bundle. Formally, it can be shown that for each fuel k, there is a threshold emission intensity, γ_k^* such that for each fuel more intensive than γ_k^*, increasing its prices will lead to a decrease in emission intensity in that industry ($\frac{\partial e_s(A)}{\partial p_k}\Big|_{\gamma_k>\gamma_k^*} < 0$) while the opposite is true for each fuel bellow γ_k^*, where

$$\gamma_k^* = p_k \frac{\sum_{\ell \neq k} \gamma_\ell (\lambda_\ell / p_\ell)^{-\sigma}}{\sum_{\ell \neq k} p_\ell^{\sigma-1} \lambda_\ell^{\sigma}} \tag{3.9}$$

For example, an increase in the relative price of fuels that pollute above γ_k^* will lead to substitution towards less polluting fuels and a decrease in emission intensity, which is exactly what happens when firms face a carbon tax.

GHG emission and carbon leakage

Now that I have characterized the production structure in this economy, I can define equilibrium GHG emissions for a firm with productivity A, which is, by definition, the product of output quantity and emission intensity:

$$\begin{aligned}
GHG_s(A) &= Y_s(A)/A && \times\, e_s(A)A \\
GHG_s(A, \text{No tax}) &= \Big(\frac{\rho-1}{\rho}\Big)\frac{1}{\tilde{p}}\Big(\frac{A}{\tilde{A}}\Big)^{\rho-1}\beta C && \times\, \frac{\sum_\ell \gamma_\ell \tilde{p}_\ell^{-\sigma}}{\tilde{p}^{-\sigma}} \\
GHG_s(A, \text{Asymmetric tax}) &= \underbrace{\Big(\frac{\rho-1}{\rho}\Big)\frac{\tilde{p}_s^{-\rho}}{(\tilde{p}_r^{1-\rho}N^r + \tilde{p}_u^{1-\rho}N^u)}\Big(\frac{A}{\tilde{A}}\Big)^{\rho-1}\beta C}_{\text{scale factor}} && \times\, \underbrace{\frac{\sum_\ell \gamma_\ell \tilde{p}_{\ell s}^{-\sigma}}{\tilde{p}_s^{-\sigma}}}_{\text{process factor}}
\end{aligned}$$

First, the scale factor is defined by macro parameters and relative perceived input price indices between regions. In contrast, the process factor is defined by relative input prices between different fuels that compose the input price index. As such, firms that are regulated face a higher

[5]For this to be valid, however, I assume that fuel taxes/prices do not affect the price of unobserved inputs like capital and labor.

relative perceived input price index and will produce less than before (scale effect). Moreover, regulated firms see a greater increase in the input price of the most polluting fuels, and they will switch towards cleaner fuels, which will decrease emission intensity (process effect). From the perspective of unregulated firms, the increase in the perceived price index of regulated firms will increase their residual demand, and they will increase output. However, unregulated firms do not see any variation in relative fuel prices, and their emission intensity remains the same. The combination of higher output but unchanged emission intensity creates carbon leakage.

3.4 Data

The empirical application of this model will be a study of the BC carbon tax. I am initially considering three main fuel types: Natural Gas, Oil, and Coal. Since this is a closed-economy model, I only look at competition between Canadian manufacturing firms.

3.4.1 Pollution and Fuel Data

The primary dataset used for plant-level pollution is the National Pollutant Inventory Release (NPRI), which records each pollutant emitted by most Canadian plants since 2000. This dataset records 280 pollutants, but I will only look at five of the most relevant pollutants to differentiate different fuels (see the identification section for more details): Carbon Monoxide (CO), Mercury (H_g), Sulphur Dioxide (SO_2), Nitrogen Oxides (NO_x) and Particulate Matters (PM). Keeping only manufacturing plants, I get between 700 and 900 plants annually between 2002 and 2015.[6] To recover fuel quantities in mmBtu at the plant level, I invert each fuel's chemical reaction under standard stationary combustion practices[7]. Thus, I get five equations (pollutants) with three unknowns (fuels) that I solve by least squares minimization subject to a non-negativity constraint, and I do this procedure for each plant each year.

[6] I chose this time frame because most firms did not report to the NPRI before 2002 and because many other Canadian provinces introduced carbon taxes and other environmental regulations in 2016.

[7] There can be heterogeneity in the quantity of each pollutant released by each fuel primarily due to combustion efficiency, which will contribute to measurement error. However, heterogeneity in pollutants across fuels tends to be much larger.

$$
\begin{aligned}
CO_{it} &= (Coal_{it} \times CO_c) + (Natgas_{it} \times CO_c) + (Oil_{it} \times CO_o) \\
H_{g,it} &= (Coal_{it} \times H_{g,c}) + (Natgas_{it} \times H_{g,ng}) + (Oil_{it} \times H_{g,o}) \\
SO_{2,it} &= (Coal_{it} \times SO_{2,c}) + (Natgas_{it} \times SO_{2,ng}) + (Oil_{it} \times SO_{2,o}) \\
NO_{x,it} &= (Coal_{it} \times NO_{x,c}) + (Natgas_{it} \times NO_{x,ng}) + (Oil_{it} \times NO_{x,o}) \\
PM_{it} &= (Coal_{it} \times PM_c) + (Natgas_{it} \times PM_c) + (Oil_{it} \times PM_o)
\end{aligned}
$$

Below is a table created from information provided by the U.S. Environmental Protection Agency (EPA) of the mapping between one mmBtu of each fuel and quantities (in pounds) of each pollutant released in the atmosphere:

Pollutant (mmBtu/lb)	Natual Gas	Oil	Coal
Carbon monoxide (CO)	0.04	0.033	0.208
Mercury (H_g)	0.001	1.122	2.591
Sulfur dioxide (SO_2)	0092	0.448	0.457
Nitrogen oxides (NO_x)	0.007	0.084	2.744
Particulate Matters (PM)	0	0.000007	0.000016

Table 3.1: Comparison of Emission Factors From Natural Gas, Oil, and Coal Combustion

Source: EPA — AP-42 (2024). Notes: This table is constructed by averaging different measures of stationary emissions from the EPA's Compilation of Stationary Emission Factors (AP-42). Natural gas is a more homogeneous fuel than coal and oil. The emissions factors for coal are created by averaging emissions for different coal grades (anthracite, bituminous, and subbituminous). Similarly, the emissions factors for oil are the average emission factors for different types of fuel oils used in manufacturing, referring to different types of distillate and residual oils.

To validate this procedure, I compare an aggregate of estimated fuel shares from each plant to aggregate fuel shares available in the following Statistics Canada database: Manufacturing industries, total annual energy fuel consumption in gigajoules (Table: 25-10-0025-01). I also cross-checked this procedure with another dataset of U.S. firms that contained both plant-level fuel quantities and pollutants.

My estimation procedures recover oil shares quite well, but coal shares are underestimated relative to natural gas shares. This is partly because only a few firms use coal in Canada, with aggregate shares nearing 1%, potentially leading to unstable estimates.

Greenhouse gases

Lastly, these pollutants are not greenhouse gases contributing to climate change (e.g., carbon dioxide CO_2, methane CH_4, and nitrous oxide N_2O) because the NPRI does not report green-

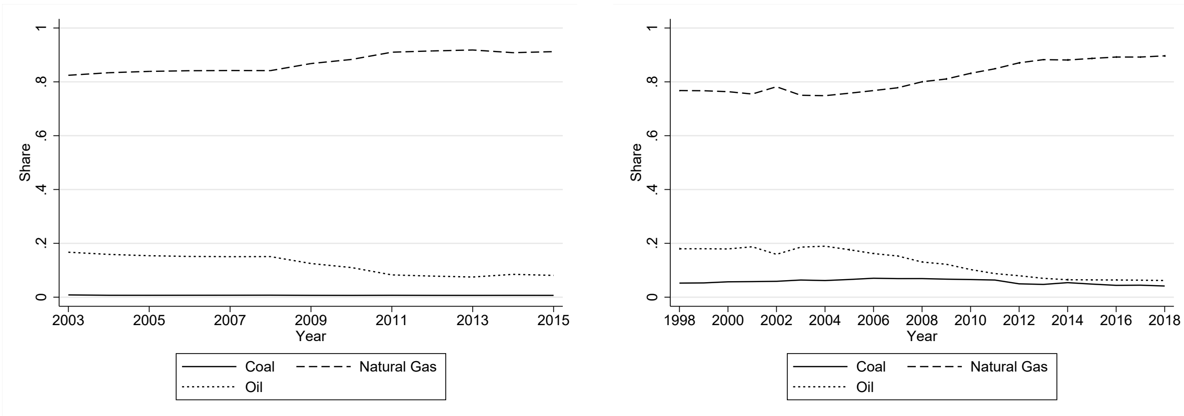

(a) Aggregated From Plant-Level Estimates (b) Statistics Canada

Figure 3.2: Aggregate fuel shares

house gases. Once I back out fuel quantities, I then use another conversion table from the EPA to convert fuel quantities to greenhouse gases. Using the GHG emission factors from Table 3.1, I then create a single measure of carbon dioxide equivalent CO_{2e} using the global warming potential method (EPA — GHG Emission Factors, 2024). This is the final measure I use in the paper to construct emissions.

	Carbon dioxide kg CO_2/mmBtu	Methane g CH_4/mmBtu	Nitrous Oxide g N_2O/mmBtu	Carbon dioxide equivalent kg CO_{2e}/mmBtu (100-Year GWP)
Coal	98.2	11	1.6	98.9
Oil	70.5	3	0.6	70.7
Natural Gas	56.1	1	0.1	56.2

Table 3.2: GHG Emission Factors

Source: EPA — GHG Emission Factors (2024). Notes: The last column is the final measure of GHG emissions.

While I use this procedure for only three fuels and five pollutants, it can be extended to include many more fuels and pollutants to improve accuracy. It can also be extended to study fuel usage by geolocated plants from remote sensing data. This approach may be interesting in regions that lack the regulatory body necessary to compile a dataset such as the NPRI.

3.4.2 Other Data

Fuel price data

Although there may be regional differences in fuel prices depending on the plant's location, I only look at global fuel prices to ensure the validity of the homogeneity assumption. Global fuel prices come from the World Bank database and are denominated in different units. To get around this issue, researchers usually convert fuel quantities into million British thermal units (mmBtu), which will be the unit of measure going forward. I use an average of crude oil prices across different sources of oil. For natural gas, I use the average price in the U.S.; for coal, I use the price of Australian coal as it is the most commonly used. Since prices are in USD, I convert them to nominal CAD prices.

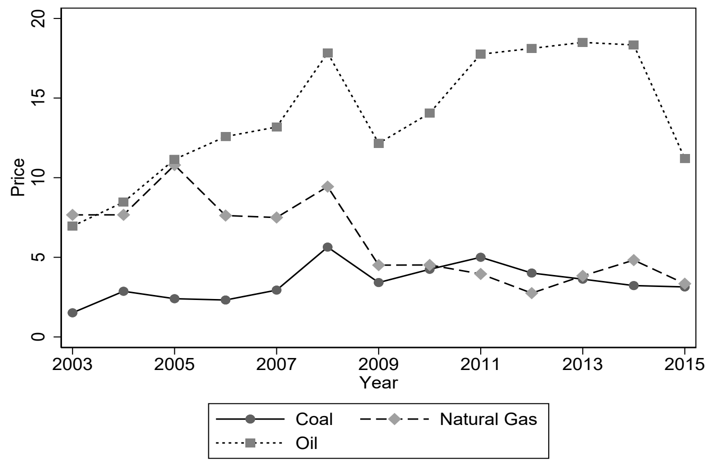

Figure 3.3: Fuel Prices in Unregulated Provinces ($CAD/mmBtu)

Carbon tax

To construct the fuel-specific tax rates, I map one mmBtu of each fuel in tons of CO_{2e} equivalent using γ_ℓ which is calculated in Table 3.2, and I multiply this by the level of the carbon tax, which started at 10$/*ton* in 2008 and increased by 5$/*ton* every year until reaching 35$/*ton* in 2013. Additionally, a much smaller carbon tax was introduced in 2008 in Quebec at 3.5$/*ton*:

	British Columbia						Quebec
	2008	2009	2010	2011	2012	2013-2015	2008-2015
Natural Gas ($/mmBtu)	0.53	0.8	1.06	1.33	1.59	1.86	0.19
Oil ($/mmBtu)	0.72	1.08	1.43	1.79	2.15	2.51	0.25
Coal ($/mmBtu)	0.99	1.48	1.98	2.47	2.97	3.44	0.35
CO_{2e} ($/ton)	10	15	20	25	30	35	3.5

Table 3.3: Carbon Tax Rates

Aggregate data

To get aggregate consumption over time, C_t, I use the share of Canadian nominal GDP from manufacturing. Then, I can get β_{jt}, the shares of Canadian manufacturing GDP that come from each 3-digit NAICS manufacturing industry. Since GDP is in dollar amount, this is going to be the share of manufacturing consumption associated with each industry: $\beta_{jt} = \frac{Y_{jt}P_{jt}}{C_j}$. Through period-by-period variation in C_t and β_{jt}, I allow for industry-specific demand shocks. To see this, if there is a one-time positive demand shock to industry j, this will increase aggregate demand C_t while simultaneously increasing β_{jt} and decreasing β_{kt} $\forall k \neq j$ such that the amount of consumption going to all other industries remains unchanged and $\sum_j \beta_{jt} = 1$ is still satisfied. Additionally, since there is only one big market in the model, a region-specific demand shock would effectively affect all firms regardless of their location. It would be captured by variation in C_t without changing relative industry shares.

3.5 Identification

Since the empirical application of this model is the B.C. and Quebec carbon taxes, I am only considering the universe of Canadian manufacturing firms under the plausible assumption that those firms take input prices as given and cannot affect global input prices with their individual decisions. As such, there will be no variation in aggregate fuel demand induced by carbon regulation that would affect equilibrium fuel prices, commonly referred to as the fossil-fuel channel of carbon leakage in the literature (Fowlie and Reguant, 2018). This means that I can exploit observed fuel price variation to identify parameters in the model. In the data, such fuel price variation is due to global shocks such as the fracking boom that massively decreased natural gas prices during the sample period and variation in the supply of oil-producing countries. Added to these global shocks is the carbon tax that creates spatial variation in fuel prices and will be critical to identifying the elasticity of substitution across firms.

There are two sets of parameters in the model: parameters related to the production technology, such as baseline fuel shares, and interfuel substitutability: $\{\{\lambda_{\ell j}\}_{=1}^{L}\}_{j=1}^{L}, \sigma\}$ and parameters about firm-level output decisions (macro parameters), such as the elasticity of demand and the distribution of productivity: $\{\rho, \sigma_A^2\}$ and parameters of the productivity distribution. Ideally, if I had data on firm-level output and emission intensity, I could separately identify both sets of parameters like Shapiro and Walker (2018), Ganapati et al. (2020), and Aichele and Felbermayr (2015) since technology parameters only appear in emission intensity and other parameters only appear in firm-level output.

However, I only observe plant-level pollutants released along with aggregate industry statistics. An important contribution of this paper is to show how to identify the direct and leakage effects of a counterfactual carbon policy on emissions under such flexible data requirements. To see this, I show in the following section how I can separate the estimation into two stages. The first stage relies on *relative* fuel quantities within plants to identify both technology and interfuel substitution parameters in a closed form using a linear, seemingly unrelated regression model (SUR). In the second stage, I use the fuel quantities in *level* net of a fuel-specific component to measure the common scale component across fuels. This common scale component contains information about firm productivity and industry aggregates, which I use to identify macro parameters, such as the distribution of productivity and the elasticity of demand.

Before showing identification results, it is also to show some key equilibrium equations when I reintroduce the industry subscript j and year subscript t into the model:

$$\text{Composite fuel (scale): } F_{ijst} = \frac{\rho-1}{\rho}\frac{\tilde{p}_{jst}^{-\rho}}{(\tilde{p}_{jrt}^{1-\rho}N_{jt}^{r}+\tilde{p}_{jut}^{1-\rho}N_{jt}^{u})}\beta_{jt}C_t\Big(\frac{A_{jst}}{\tilde{A}_{jt}}\Big)^{\rho-1}$$

$$\text{Individual fuel quantities: } q_{ijst}^{\ell} = \Big(\frac{\tilde{p}_{\ell jst}}{\tilde{p}_{jst}}\Big)^{-\sigma}F_{ijst}$$

3.5.1 Stage 1: Identification of Technology Parameters

With estimated fuel quantities, I can use optimality conditions from the model to estimate interfuel substitution and fuel efficiency parameters in a seemingly unrelated regression model. Indeed, each fuel contains a fuel-specific factor driven by its "perceived price" and a common factor to all fuels that is proportional to the output of a given firm:

$$\begin{aligned} q_{ijst}^{\ell} &= \Big(\frac{\tilde{p}_{\ell jst}}{\tilde{p}_{jst}}\Big)^{-\sigma}F_{ijst} \\ &= \tilde{p}_{\ell jst}^{-\sigma}X_{ijst} \\ &= \underbrace{\Big(\frac{p_{\ell st}}{\lambda_{\ell j}}\Big)^{-\sigma}}_{\text{fuel-specific}}\ \underbrace{X_{ijst}}_{\text{common scale}} \end{aligned}$$

Where F_{ijst} is the factor common to all fuels. Since there are more than three fuels in reality, and since the pollutant factor to fuel conversion approximates chemical reactions under average conditions, it is likely that there is some measurement error in implied fuel quantities, $\exp(u_{ijst}^{\ell})$. Measurement error comes from unobservables such as heat at which combustion happens, the type of coal or oil product used, and the technology used for combustion, which I assume to be iid across time and firms:

$$\hat{q}^{\ell}_{ijst} = \left(\frac{p_{\ell st}}{\lambda_{\ell ij}}\right)^{-\sigma} X_{ijst} \exp(u^{\ell}_{ijst})$$

Taking logs,

$$\ln q^{\ell}_{ijst} = -\sigma \ln p_{\ell st} + \sigma \ln \lambda_{\ell j} + \ln X_{ijst} + u^{\ell}_{ijst}$$

Since all fuels contain the common factor, I can eliminate it by subtracting each fuel $\ell = 2, 3$ by the first fuel, which is what allows me to identify technology parameters separately form the rest.

$$\begin{aligned}
\ln q^2_{ijst} - \ln q^1_{ijst} &= \sigma(\ln p_{1st} - \ln p_{2st}) + \sigma(\ln \lambda_{j2} - \ln \lambda_{j1}) + (u^2_{ijst} - u^1_{ijst}) \\
\ln q^3_{ijst} - \ln q^1_{ijst} &= \sigma(\ln p_{1st} - \ln p_{3st}) + \sigma(\ln \lambda_{j3} - \ln \lambda_{j1}) + (u^3_{ijst} - u^1_{ijst})
\end{aligned}$$

Intuitively, if the price of the first fuel increases, firms will switch towards fuels 2 and 3, which will depend on the elasticity of substitution σ. I can rewrite this system of equations in matrix form:

$$\begin{bmatrix} \ln q^2_{ijst} - \ln q^1_{ijst} \\ \ln q^3_{ijst} - \ln q^1_{ijst} \end{bmatrix} = \sigma \begin{bmatrix} \ln p_{1st} - \ln p_{2st} \\ \ln p_{1st} - \ln p_{3st} \end{bmatrix} + \begin{pmatrix} \Gamma_{2j} \\ \Gamma_{3j} \end{pmatrix} + \begin{bmatrix} \epsilon^2_{ijst} \\ \epsilon^3_{ijst} \end{bmatrix}$$
$$\mathbf{\Delta} \ln \mathbf{q_{ijst}} = \sigma \mathbf{\Delta} \ln \mathbf{p_{st}} + \mathbf{\Gamma_j} + \boldsymbol{\epsilon_{ijst}}$$

Where $\Gamma_{\ell j} = \sigma(\ln \lambda_{j\ell} - \ln \lambda_{j1})$ and $\epsilon^{\ell}_{ijst} = u^{\ell}_{ijst} - u^1_{ijst}$. this is a SUR model because ϵ^2_{ijst} is correlated with ϵ^3_{ijst} but uncorrelated across firms and across time. Identification of the elasticity of substitution comes from variations in aggregate fuel prices over time, along with the carbon taxes that were implemented in B.C. and Quebec. Estimating $\{\sigma, \Gamma_{2j}, \Gamma_{3j}\}$ is a reduced form that will recover all technology parameters as follows:

$$\exp(\Gamma_{2j}/\sigma) = \frac{\lambda_{2j}}{\lambda_{1j}}$$
$$\exp(\Gamma_{3j}/\sigma) = \frac{\lambda_{3j}}{\lambda_{1j}}$$
$$\lambda_{1j} = 1 - \lambda_{2j} - \lambda_{3j]}$$

$\lambda_{1j}, \lambda_{2j}, \lambda_{3j}$ will then be the solution to the above system of three equations. Since I am only keeping oil and natural gas in the model, the fuel-specific efficiency terms can be recovered in closed form:

$$\hat{\lambda}_{ng,j} = \frac{\exp(\hat{\Gamma}_j/\hat{\sigma})}{\exp(\hat{\Gamma}_j/\hat{\sigma}) + 1}$$
$$\hat{\lambda}_{oil,j} = \frac{1}{\exp(\hat{\Gamma}_j/\hat{\sigma}) + 1}$$

3.5.2 Stage 2: Identification of Macro Parameters

Upon estimating technology parameters, I can use the levels of log fuel quantities subtracted from a substitution term to recover the common scale component, F_{ijst} plus fuel-specific measurement error:[8]

$$\ln q^{\ell}_{ijst} + \hat{\sigma}(\ln \tilde{p}_{\ell jst} - \ln \tilde{p}_{jst}) = \ln F_{ijst} + u^{\ell}_{ijst} \quad \forall \ell \in \{oil, gas, coal\} \tag{3.10}$$

I can open up this expression in equilibrium because F_{ijst} is just an industry and time-augmented version of equation 3.7. It is a function of the elasticity of demand, ρ, unobserved aggregate productivity $\tilde{A}_{jt}$, unobserved firm-specific productivity A_{ijst}, unobserved fuels-specific measurement error u^{ℓ}_{ijst}, and many quantities that are now observed (input price indices $\tilde{p}_{jst}$, share of firms in different regions $\tilde{N}^{s}_{jt}$, industry output share β_{jt}, and aggregate output C_t). Hence, in matrix form, we can write this as a system of estimating equations:

[8] Note that F_{ijst} differs from the "true" scale component X_{ijst} because I also subtract input price index in equation 3.10, which is common to all fuels.

$$\underbrace{\ln \mathbf{q}^{\ell}_{\mathbf{ijst}} + \hat{\sigma}(\ln \tilde{\mathbf{p}}_{\ell\mathbf{jst}} - \ln \tilde{p}_{jst})}_{y_{ijst}} = \underbrace{\ln\Big(\frac{\rho-1}{\rho}\Big) + \ln\beta_{jt} + \ln C_t - \rho\ln\tilde{p}_{jst} - \ln(\tilde{p}^{1-\rho}_{jrt}\tilde{N}^{r}_{jt} + \tilde{p}^{1-\rho}_{jut}\tilde{N}^{u}_{jt})}_{X_{ijst}(\rho)}$$
$$+ \underbrace{(\rho-1)(\ln A_{ijst} - \ln\tilde{A}_{jt}) + \mathbf{u}^{\ell}_{\mathbf{ijst}}}_{\epsilon_{ijst}} \quad (3.11)$$

Where aggregate productivity in each industry was defined as:

$$\tilde{A}_{jt} = \Big(\int_0^{\infty} A^{\rho-1} g_{jt}(A)dA\Big)^{1/(\rho-1)}$$

I now assume a distribution for productivity and fuel-specific measurement error, allowing me to estimate this system with maximum likelihood. I assume that productivity is distributed log-normal with mean and variance common to all firms in all years $\tilde{A}_{jt} = \tilde{A}$, and that measurement error is normally distributed with mean zero and fuel-specific variance. I also allow for measurement error to be correlated across fuels. The marginal distributions are then:

$$\begin{aligned} &A_{ijst} \sim LN(\mu, \sigma^2_A) \iff \ln A_{ijst} \sim N(\mu, \sigma^2_A) \\ &u^{\ell}_{ijst} \sim N(0, \sigma^2_{\ell}) \end{aligned}$$

Using the moment generating function for a log-normal distribution, I can get a closed-form solution for $\tilde{A}_{jt} = \tilde{A}$ (see appendix for details):

$$\tilde{A} = \Big(\exp\Big(\mu(\rho-1) + (\rho-1)^2\sigma^2_A/2\Big)\Big)^{1/(\rho-1)} \quad (3.12)$$

With that in mind, the relative log productivity term will be normally distributed, such that the vector of fuel-specific error terms in estimating equation 3.11, ϵ_{ijst}, will be jointly normal (see appendix for details):

$$(\rho - 1)(\ln A_{ijst} - \ln \tilde{A}) \sim N\Big(-(\rho - 1)^2\sigma_A^2/2, (\rho - 1)^2\sigma_A^2\Big)$$

$$\begin{pmatrix} \epsilon_{ijst}^{ng} \\ \epsilon_{ijst}^{o} \end{pmatrix} \sim N\left(\begin{pmatrix} -(\rho - 1)^2\sigma_A^2/2 \\ -(\rho - 1)^2\sigma_A^2/2 \end{pmatrix}, \begin{pmatrix} (\rho - 1)^2\sigma_A^2 + \sigma_{ng}^2 & (\rho - 1)^2\sigma_A^2 + \sigma_{ng,o} \\ (\rho - 1)^2\sigma_A^2 + \sigma_{ng,o} & (\rho - 1)^2\sigma_A^2 + \sigma_o^2 \end{pmatrix} \right)$$

From this, I can construct the likelihood as the product of n bivariate normal density, take the natural logarithm, and estimate all the parameters: $\big(\rho, \sigma_{ng}^2, \sigma_o^2, \sigma_A^2, \sigma_{ng,o}\big)$.

To get a bit of intuition behind estimating the within-industry elasticity of substitution across firms ρ, note that input price variation only identifies ρ when it is spatially differentiated, in this case through the asymmetric carbon tax. To see this, if input prices are the same everywhere in a given period ($\tilde{p}_{jrt} = \tilde{p}_{jurt} = \tilde{p}_{jt}$), then firms' choice of the composite fuel index collapses to:

$$\ln F_{ijst} = \ln\Big(\frac{\rho - 1}{\rho}\Big) + \ln\beta_{jt} + \ln C_t - \ln \tilde{p}_{jt}$$

Intuitively, this happens because variation in input prices is the same for all firms in the industry, hence there is no incentive for consumers to substitute across firms. With the introduction of a tax, consumers will switch towards unregulated firms at rate $\rho > 1$ because the tax raises the marginal costs of regulated firms, which raises their output prices due to the monopolistic competition pricing rule.

3.6 Results

Technology parameters

Below are estimates of technology parameters. Since quantities of estimated coal are very low and only represent 1% of fuel consumption, I only keep natural gas and oil in the analysis. The standard error for the interfuel substitution parameter $\hat{\sigma}$ is an OLS standard error, and the standard errors for the fuel shares parameters are derived using the Delta Method (see appendix for details). I also construct 95% bootstrap confidence intervals.

Many firms often consume a dominant fuel; the model captures these quasi-corner solutions

	Baseline Natural Gas share: $\hat{\lambda}_{ng,j}$								
Industry	311	312	314	315	321	322	323	324	325
Estimate	0.984	0.871	0.007	0.000	1.000	1.000	1.000	0.983	1.000
S.E.	(0.023)	(0.176)	(0.104)	(0.001)	(0.000)	(0.002)	(0.000)	(0.022)	(0.001)
95 % CI lb	0.931	0.444	0.000	0.000	1.000	0.995	1.000	0.936	0.996
95 % CI ub	0.999	0.993	0.224	0.000	1.000	1.000	1.000	0.999	1.000
$\sum_{t=2003}^{2015} N_{jt}$	1,005	145	40	16	2,669	1,305	24	1,036	1,378
	Baseline oil share $\hat{\lambda}_{oil,j}$								
Industry	311	312	314	315	321	322	323	324	325
Estimate	0.016	0.129	0.993	1.000	0.000	0.000	0.000	0.017	0.000
S.E.	(0.023)	(0.176)	(0.104)	(0.001)	(0.000)	(0.002)	(0.000)	(0.022)	(0.001)
95 % CI lb	0.001	0.007	0.776	1.000	0.000	0.000	0.000	0.001	0.000
95 % CI ub	0.069	0.556	1.000	1.000	0.000	0.005	0.000	0.064	0.004
$\sum_{t=2003}^{2015} N_{jt}$	1,005	145	40	16	2,669	1,305	24	1,036	1,378
	Baseline Natural Gas share: $\hat{\lambda}_{ng,j}$								
Industry	326	327	331	332	333	335	336	337	339
Estimate	0.960	0.994	0.997	1.000	1.000	0.821	1.000	1.000	0.519
S.E.	(0.068)	(0.011)	(0.008)	(0.001)	(0.000)	(0.314)	(0.000)	(0.002)	(0.301)
95 % CI lb	0.792	0.967	0.979	0.998	1.000	0.069	1.000	0.995	0.050
95 % CI ub	0.998	1.000	1.000	1.000	1.000	0.998	1.000	1.000	0.961
$\sum_{t=2003}^{2015} N_{jt}$	195	903	980	189	57	42	328	99	121
	Baseline oil share $\hat{\lambda}_{oil,j}$								
Industry	326	327	331	332	333	335	336	337	339
Estimate	0.040	0.006	0.003	0.000	0.000	0.179	0.000	0.000	0.481
S.E.	(0.068)	(0.011)	(0.008)	(0.001)	(0.000)	(0.314)	(0.000)	(0.002)	(0.301)
95 % CI lb	0.002	0.000	0.000	0.000	0.000	0.002	0.000	0.000	0.039
95 % CI ub	0.208	0.033	0.021	0.002	0.000	0.931	0.000	0.005	0.950
$\sum_{t=2003}^{2015} N_{jt}$	195	903	980	189	57	42	328	99	121
	Interfuel elasticity of substitution: $\hat{\sigma}$								
Estimate	1.856								
S.E.	0.416								
95 % CI lb	1.040								
95 % CI ub	2.671								
N_{total}	10,532								

Table 3.4: Estimates of Technology Parameters

by $\hat{\lambda}_{j,\ell}$ close to either one or zero. Since natural gas is by far the most prominent fuel, $\hat{\lambda}_{j,oil}$ is often close to zero, which leads to a perceived price distribution that is much larger for oil than natural gas such that firms in most industries would never want to purchase oil:

	N	Mean	Std. Deviation	Min	Max
$\tilde{p}_{oil}$	10,532	1.62e+08	2.86e+08	6.957819	9.32e+08
$\tilde{p}_{natgas}$	10,532	645.9813	17690.24	2.752033	797652.4

Table 3.5: Summary statistics, perceived prices ($\tilde{p}_{\ell jt} = \frac{p_{\ell t}}{\lambda_{j\ell}}$)

Macro parameters

To get standard errors, I bootstrapped the entire dataset by blocks of individual firms, each sample being every year a firm operates. This allows for within-firm correlation across years, which is to be expected.

	Within-industry elasticity of substitution across firms	Variance of productivity	Variance of measurement error, natgas	Variance of measurement error, oil	Covariance between measurement error
	$\hat{\rho}$	$\hat{\sigma}^2_A$	$\hat{\sigma}^2_{ng}$	$\hat{\sigma}^2_o$	$\hat{\sigma}_{ng,o}$
Estimate	2.242	1.623	71.5	188.184	-39.86
S.E.	(0.7837)	(1.909)	(2.77)	(4.68)	(3.28)
95 % C.I.	[1.155, 3.57]	[0.46,6.27]	[67.72,77.79]	[182.11,196.56]	[-44.16,33.98]
Observations	10,532	10,532	10,532	10,532	10,532

Table 3.6: Estimates of macro parameters

3.7 Counterfactuals

3.7.1 Causal Effects of the B.C. And Quebec Carbon Taxes

I now have everything needed to recover counterfactual GHG emissions levels had the carbon tax not been introduced in B.C. and Quebec. In this counterfactual, regulated firms are always subject to unregulated input prices, and firms in other provinces do not get competitive gains from the input taxes raising input prices in B.C. and Quebec. Doing this exercise for regulated and unregulated firms separately will recover the tax's direct effect and the tax's carbon leakage effect separately. Let T_τ be the year the tax was implemented. GHG emissions then refer to the sum of GHG across firms for a given industry and year post-treatment (inner sum), the sum over all industries in a given year (middle sum), and the sum over all years post-treatment (outer sum). I do this for regulated and unregulated firms separately. I compare total GHG emissions generated by the model when prices include the tax (in red) subtracted to GHG emissions when prices exclude the tax (in blue) to get the effect of the carbon taxes.

$$\text{Direct effect (BC firms): } \left(\sum_{t\geq T_\tau}\sum_j\sum_{i:r}\widehat{GHG}_{ijt}(\tilde{p}_{jrt},\tilde{p}_{jut})\right)-\left(\sum_{t\geq T_\tau}\sum_j\sum_{i:r}\widehat{GHG}_{ijt}(\tilde{p}_{jrt}=\tilde{p}_{jut},\tilde{p}_{jut})\right)=$$

$$=\sum_{t\geq T_\tau}\sum_j\sum_{i:r}\frac{\hat{\rho}-1}{\hat{\rho}}\beta_{jt}C_t\Big(\frac{A_{ijrt}}{\tilde{A}}\Big)^{\hat{\rho}-1}\left(\underbrace{\frac{\tilde{p}_{jrt}^{-\hat{\rho}}}{(\tilde{p}_{jrt}^{1-\hat{\rho}}+\tilde{p}_{jut}^{1-\hat{\rho}})}}_{\text{scale channel}}\underbrace{\frac{\sum_\ell\gamma_\ell\tilde{p}_{\ell jrt}^{-\hat{\sigma}}}{\tilde{p}_{jrt}^{-\hat{\sigma}}}}_{\text{process channel}}-\underbrace{\frac{1}{\tilde{p}_{jut}}}_{\text{scale channel}}\underbrace{\frac{\sum_\ell\gamma_\ell\tilde{p}_{\ell jut}^{-\hat{\sigma}}}{\tilde{p}_{jut}^{-\hat{\sigma}}}}_{\text{process channel}}\right)$$

$$\text{Carbon leakage (Other firms): } \left(\sum_{t\geq T_\tau}\sum_j\sum_{i:u}\widehat{GHG}_{ijt}(\tilde{p}_{jrt},\tilde{p}_{jut})\right)-\left(\sum_{t\geq T_\tau}\sum_j\sum_{i:u}\widehat{GHG}_{ijt}(\tilde{p}_{jrt}=\tilde{p}_{jut},\tilde{p}_{jut})\right)=$$

$$=\sum_{t\geq T_\tau}\sum_j\sum_{i:u}\frac{\hat{\rho}-1}{\hat{\rho}}\beta_{jt}C_t\Big(\frac{A_{ijut}}{\tilde{A}}\Big)^{\hat{\rho}-1}\frac{\sum_\ell\gamma_\ell\tilde{p}_{\ell jut}^{-\hat{\sigma}}}{\tilde{p}_{jut}^{-\hat{\sigma}}}\underbrace{\left(\frac{\tilde{p}_{jrt}^{-\hat{\rho}}}{(\tilde{p}_{jrt}^{1-\hat{\rho}}+\tilde{p}_{jut}^{1-\hat{\rho}})}-\frac{1}{\tilde{p}_{jurt}}\right)}_{\text{scale channel}}$$

Moreover, I can decompose both the direct and carbon leakage effects into a scale channel (change in aggregate output) and a process/technique channel (change in aggregate emission intensity due to interfuel substitution). There is no process effect for unregulated firms because they do not face the carbon tax, and relative fuel prices stay constant. I also did this exercise for Quebec firms that faced a smaller carbon tax around the same period. Hence, the effect of the carbon tax on unregulated firms will be calculated by setting both the Quebec and B.C. carbon taxes to zero. I estimate an overall decrease in GHG emissions of 37.6 % in BC and 4.87 % in Quebec and an increase in GHG emissions of 2.6 % in unregulated provinces, which can be attributed to carbon leakage.

While aggregate percentage changes in emissions are suggestive that emissions decreased much more in British Columbia than they increased in unregulated provinces, this is misleading because the level of emissions in unregulated provinces is much larger at the baseline. There are many more firms in unregulated provinces, and those firms tend to concentrate on more polluting industries. As a result, significant leakage mitigates 45% of emissions reduction efforts.

In the next section, I document which effect drives most of this variation in GHG emissions. I also show what happens under a uniform carbon tax across all provinces, which the federal government introduced in 2018 with the Greenhouse Gas Pollution Pricing Act.

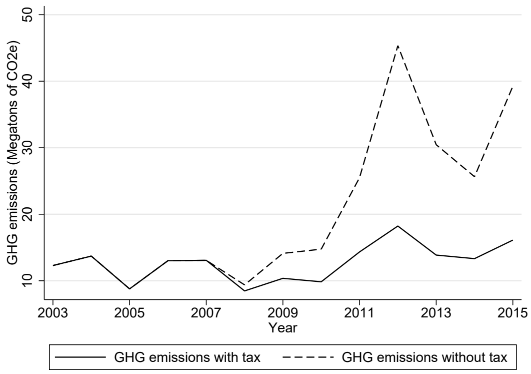

(a) British Columbia

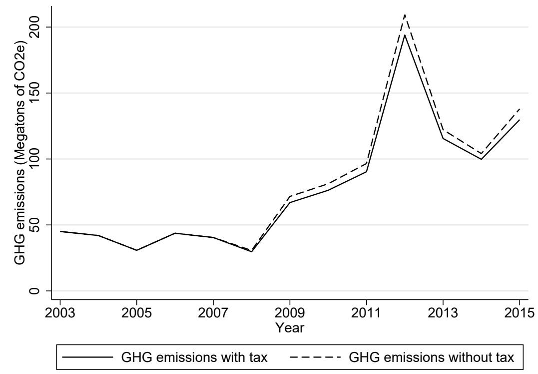

(b) Quebec

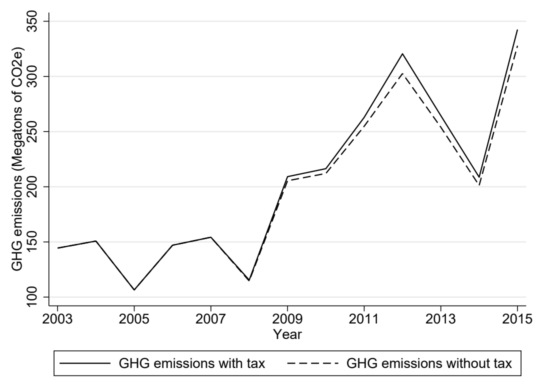

(c) Unregulated Provinces

Figure 3.4: Estimated and Counterfactual GHG Emissions (Kilotons CO_{2e})

	Aggregate Emissions (kilotons CO_{2e}) (2003-2015) **No Tax**	Aggregate Emissions (kilotons CO_{2e}) (2003-2015) **Carbon Tax**	Level Difference (Carbon Tax - No Tax)	Percentage Difference (Carbon Tax - No Tax)
British Columbia	265	165	-100	-37.6%
Quebec	1,055	1,004	-51.4	-4.87%
Unregulated Provinces	2,575	2,643	67.4	2.6%

Table 3.7: Effect of Carbon Taxes on Emissions of Regulated and Unregulated Firms

Role of Fuel Substitution

The introduction of multiple fuels allows firms to substitute across fuels, which determines the GHG emissions process channel. Indeed, absent fuel substitution, firms only choose how much of the composite fuel they want to purchase, forming their output (scale) decisions. By including the tax only in the perceived price of composite fuel $\tilde{p}_{sjt}$ and keeping the fuel-specific perceived prices without the tax $\tilde{p}_{\ell jt}$ (and vice versa), I am able to decompose the overall effect of the carbon tax into the scaling factor and the fuel substitution factor.

My results suggest that almost all changes in GHG emissions (98.51% for B.C., 96.1% in Quebec, and 100% in unregulated provinces) are due to firms adjusting output, and only a tiny fraction is due to fuel substitution. In unregulated provinces, this is expected because unregulated firms did not face the tax directly, so there is no scope for fuel substitution. However, this is the case in B.C. and Quebec because most firms are consuming a dominant fuel, with the perceived price of other fuels being extremely high. Most dominant fuels tend to be natural gas, and the carbon tax raised the price of oil more than it raised the price of natural gas because the former releases more $CO_{2e}/mmBtu$. Thus, there is a limited role for substituting towards the newly relatively cheaper fuel (natural gas) because most firms already use large quantities of natural gas. The lack of a cleaner alternative to natural gas partly explains why carbon leakage is so large.

	Effect of carbon tax on GHG emissions (%)		
	British Columbia	Quebec	Unregulated Provinces
Scale	-37.1	-4.68	2.62
Fuel substitution	-1.2	-0.2	0
Scale and fuel substitution	-37.66	-4.87	2.62

Table 3.8: Difference in GHG Emissions With Counterfactual No Policy, Percentage

3.7.2 Uniform Policy

An interesting aspect of this model is that I can impose a uniform carbon tax on all firms in all provinces such that there is no risk of carbon leakage, and I can see how it compares with the actual policy. I find that the reduction in GHG emissions would be between 19% and 22% depending on the province, which is lower than the 37.66% decrease of the actual policy in British Columbia

Interestingly, this happens because as all firms face the same tax rate, there is no competitive hedge from being in one province versus another. All the production that would have shifted to unregulated provinces due to the asymmetry of the carbon tax stays in the province of origin.

	Change in GHG emissions (%)
British Columbia	-23.77
Quebec	-22.82
Unregulated Provinces	-19.90

Table 3.9: Effect on GHG Emissions of Introducing a Uniform Carbon Tax Across Canada With a Tax Rate Equivalent to the B.C. Carbon Tax (Percentage)

The aggregate emission decrease is simply due to a uniform increase in marginal costs. Overall, a uniform carbon tax is much better at reducing aggregate GHG emissions, estimated at 21% relative to a total decrease of 2.15% in the case of the asymmetric tax.

	Total change in GHG emissions (%)
BC carbon tax	-2.15
Uniform carbon tax	-21.15

Table 3.10: Total Effect of Implementing a Uniform Carbon Tax Across Canada Relative to Total Effect of B.C. Carbon Tax (Percentage)

3.8 Conclusion

In this paper, I build a production model where manufacturing firms can substitute between fuels and compete across different regions. Using a panel of publicly available emissions data from Canadian plants, I can recover counterfactual emissions of regulated and unregulated plants following British and Quebec carbon taxes implemented in 2007 and 2008, respectively. I find that carbon leakage increased GHG emissions by 2.62 % in unregulated provinces relative to a decrease in emissions by 37.6% in British Columbia and 4.87% in Quebec and that most of the decrease in emissions by regulated plants was due to variation in output rather than substitution towards cleaner fuels.

However, these percentage differences hide a much larger effect of carbon leakage when looking at the level of emissions. Indeed, both carbon taxes accounted for a decrease in emissions by 150 kilotons of CO_{2e} by regulated firms relative to an increase in emissions of 67 kilotons of CO_{2e} by unregulated firms. Carbon leakage thus mitigated 45% of emissions reduction efforts in B.C. and Quebec. Introducing a uniform carbon tax across Canada, as was legislated in the Greenhouse Gas Pollution Pricing Act of 2018, reduces total Canadian greenhouse gas emissions by 21 %. Relative to a net Canadian-wide emissions decrease of 2.15% if the carbon tax was only implemented in B.C. and Quebec, the uniform carbon tax significantly mitigates leakage concerns.

Bibliography

Aichele, Rachel and Gabriel Felbermayr, "Kyoto and Carbon Leakage: an Empirical Analysis of the Carbon Content of Bilateral Trade," *Review of Economics and Statistics*, 2015, *97* (1), 104–115.

Atkinson, Scott E. and Rong Luo, "Estimation of Production Technologies with Output and Environmental Constraints," *International Economic Review*, 2023, *0* (0).

Böhringer, Christoph, Jared C Carbone, and Thomas F Rutherford, "The Strategic Value of Carbon Tariffs," *American Economic Journal: Economic Policy*, 2016, *8* (1), 28–51.

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer, "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?," *Journal of Policial Economy*, 2000, *108* (6).

Cho, Won G., Kiseok Nam, and José A. Pagan, "Economic growth and interfactor/interfuel substitution in Korea," *Energy Economics*, 2004, *26*, 31–50.

Copeland, Brian R and M Scott Taylor, "Trade , Growth , and the Environment," *Journal of Economic Literature*, 2004, *42* (1), 7–71.

Felder, Stefan and Thomas F Rutherford, "Unilateral CO2 Reductions and Carbon Leakage: The Consequences of International Trade in Oil and Basic Materials," *Journal of Environmental Economics and Management*, 1993, *25* (2), 162–176.
Environmental Production Agency

Environmental Production Agency, "AP-42: Compilation of Air Emissions Factors from Stationary Sources," 2024. Retrieved from `https://www.epa.gov/air-emissions-factors-and-quantification/ap-42-compilation-air-emissions-factors-stationary-sources`.
Environmental Production Agency

Environmental Production Agency, "GHG Emission Factors Hub," 2024. Retrieved from `https://www.epa.gov/climateleadership/ghg-emission-factors-hub`.

Fowlie, Meredith and Mar Reguant, "Challenges in the Measurement of Leakage Risk," *AEA Papers and Proceedings*, 2018, *108.*

Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker, "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes," *American Economic Journal: Applied Economics*, 2020, *12* (2), 303–342.

Garella, Paolo Giorgio and Maria Teresa Trentinaglia, "Carbon Tax, Emission Standards, and Carbon Leak Under Price Competition," *Environmental and Resource Economics*, 2019, *72* (4), 941–964.

Hassler, John, Conny Olovsson, and Michael Reiter, "Integrated Assessment in a Multi-region World with Multiple Energy Sources and Endogenous Technical Change," 2019.

Hoel, Michael, "Global environmental problems: The effects of unilateral actions taken by one country," *Journal of Environmental Economics and Management*, 1991, *20* (1), 55–70.

Holland, Stephen P, "Emissions taxes versus intensity standards: Second-best environmental policies with incomplete regulation," *Journal of Environmental Economics and Management*, 2012, *63* (3), 375–387.

Hyland, Marie and Stefanie Haller, "Firm-level Estimates of Fuel Substitution," *The Energy Journal*, 2018, *39* (6), 71–98.

Levinson, Arik and M Scott Taylor, "UNMASKING THE POLLUTION HAVEN EFFECT," *International Economic Review*, feb 2008, *49* (1), 223–254.

Ma, Hengyun, Les Oxley, John Gibson, and Bonggeun Kim, "China's energy economy: Technical change, factor demand and interfactor/interfuel substitution," *Energy Economics*, 2008, *30*, 2167–2183.

Melitz, Marc J., "The impact of trade on intra-industry reallocations and aggregate industry productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

_ **and Stephen J. Redding**, *Heterogeneous Firms and Trade*, Vol. 4 2014.

Pindyck, Robert S., "Interfuel Substitution and the Industrial Demand for Energy," *The Review of Economics and Statistics*, 1979, *61* (2), 169–179.

Shapiro, Joseph S. and Reed Walker, "Why is pollution from us manufacturing declining? the roles of environmental regulation, productivity, and trade," *American Economic Review*, 2018, *108* (12), 3814–3854.

Smith, Dominic and Sergio Ocampo, "The Evolution of U . S . Retail Concentration," 2020.

Wang, Xiaolei and Boqiang Lin, "Factor and fuel substitution in China's iron & steel industry: Evidence and policy implications," *Journal and Cleaner Production*, 2017, *141*, 751–759.

Worrell, Ernst, Lenny Bernstein, Joyashree Roy, Lynn Price, and Jochen Harnisch, "Industrial energy efficiency and climate change mitigation," *Energy Efficiency*, 2009, *2* (2), 109–123.

Chapter 4

Long-Term Contracts and Secondary Markets: Theory and Evidence from Natural Gas Pipelines

4.1 Introduction

In North America, pipelines are the dominant method for transporting natural gas. They are extremely costly to build, averaging $5 million/km with 10 to 20 years for the pipeline to break even on its investment. To share the risk associated with large sunk investment costs, pipelines must sign long-term contracts (typically ten years or more) with customers prior to the pipeline's construction, which entitles the customer to reserved capacity. Long-term contracts in the primary market can create situations of artificial constraint in which physical capacity is available, but contracting frictions impede the flow of natural gas. For these reasons, there is a secondary market unique to the U.S. where shippers can lease capacity to other shippers for any length of time. The secondary market is much less regulated than the primary market, with the aim of fostering a competitive environment between shippers.

In this project, we evaluate the efficiency of the U.S. natural gas pipeline network at intertemporally allocating pipeline capacity in response to demand and supply shocks through its secondary market. We study trade-offs associated with the secondary market and provide guidance for deregulation efforts in similar industries. We study shocks to large weather fluctuations that affect the supply and demand for energy, increasingly common with the advent of climate change (Bathiany et al., 2018, Wang, Biasutti, Byrne, Castro, Chang, Cook, Fu, MGrimm, Jaha, Hendon et al., 2021). Through the lens of the secondary market for capacity release, we contribute to a growing literature studying the resilience of energy systems and markets in response to such atmospheric conditions (Gonçalves, Costoya, Nieto and Liberato, 2024, Jasiūnas, Lund and Mikkola, 2021, Waseem and Manshadi, 2020).

Using novel daily transaction data for selected U.S. natural gas pipelines between 2005 and 2023, we find that the secondary market responds to large regional fluctuations in demand for natural gas. We narrow down our study to the extreme cold wave of February 2021. This cold wave significantly increased natural gas demand, particularly demand for residential heating and electricity generation in the Deep South. Notably, natural gas prices at hubs in Texas and southeastern California rose substantially during this period. In contrast, prices at hubs outside the cold wave's radius, such as in Northwestern California, did not rise.

We leverage this spatial variation in natural gas prices to identify control and treated pipelines. We choose Gas Transmission Northwest (GTN) as a control, which carries natural gas from Western Canada to Northwestern California. We chose the El Paso pipeline as a treatment, which carries natural gas to major hubs in Texas and Southeastern California. In a difference-in-difference event study framework similar to Miller (2023) and Freyaldenhoven, Hansen, Pérez and Shapiro (2024), we find that the cold wave caused a 50% increase in capacity released from buyers (replacers) to sellers (releasers) in the secondary market. We also find that the secondary market facilitates a reduction in search costs between replacers and releasers but fosters an anti-competitive environment where both sides of the market influence prices.

This secondary market is one of the only real examples of a Coasian market in which participants trade legal entitlement to natural gas pipeline capacity. Our paper is the first to study such a secondary market empirically. We aim to provide some evidence that generalizes to other similar heavily regulated markets with large sunk investment costs and to provide guidance for the organization of pipeline networks outside of the U.S.

To understand the aforementioned trade-offs better, we show that the vast majority of shippers that release capacity are end-consumers (power plants, local distribution companies that serve specific urban areas, and industrial plants) who have long-term contracts with pipelines and must pay the large reservation fees even if they do not utilize their capacity. The vast majority of replacers are marketers. Marketers are intermediaries that engage in arbitrage opportunities, buying gas in cheap hubs and selling gas in expensive hubs to consumers.

In this context, we provide empirical evidence for long-term relationships between market participants in the capacity release market and propose potential hypotheses for why such relationships exist. We show that over 95% of contracts signed are private arrangements made between shippers. Many of these arrangements recur between shippers that meet multiple times over the sample period, which we call long-term relationships. This is similar to the U.S. trucking industry, where 80% of the shipping contracts are signed through existing relationships between shippers (analogous to the replacement shippers) and carriers (analogous to releasing shippers), and only 20% of the contracts are signed via a spot market (Harris and Nguyen, 2023).

Considering the unique aspects of this market, search costs are the key factor influencing the formation of long-term relationships. This is attributed to the frequent requirement to secure capacity at various delivery and receipt locations, along with other customary contractual obligations such as the option for a releaser to be able to recall its capacity at any time, making each contract a highly differentiated product. Thus, when releasers do not need their capacity, they search for replacers to avoid paying reservation fees, whereas replacers search for releasers when facing an arbitrage opportunity. Using event studies, we find that contracts are much more likely to be signed between long-term partners during the cold wave. This suggests that the opportunity cost of searching is higher in periods of high demand and that long-term relationships reduce search costs.

Moreover, the secondary market facilitates the formation of long-term relationships. While contracts are mostly private arrangements, the Federal Energy Regulatory Commission (FERC) mandates that pipelines publicly post transaction details from the secondary market on public data-sharing platforms. We argue that this regulatory constraint is what facilitates the formation of long-term relationships.

However, the quasi-lack of regulation on the price charged for pipeline capacity in the secondary market fosters an anti-competitive environment.[1] To show this, we leverage variation in market concentration across different pipelines and regions, where some pipeline zones have dominant releasers, some have dominant replacers, and some have both. We find that a 1% increase in releaser market share is associated with a 0.1% increase in capacity prices. In contrast, a 1% increase in replacer market share is associated with a 0.04% decrease in capacity prices. The welfare implications of these price-setting behaviors are uncertain, and we plan to study them using a search model with two-sided market power in future research.

Related Literature

This paper contributes to three broad strands of literature on (i) supply-chain and energy market resilience, (i) the effectiveness of secondary markets, and (iii) the pipeline industry.

[1] There are some constraints on capacity prices in the secondary markets based on how long capacity is expected to be released. See Section 4.2 for details.

Chen (2023) shows that mergers among those networks result in significant cost reductions and a notable rise in markups. Nonetheless, there is still a cost pass-through to customers, leading to a reduction in shipment prices. Based on the findings of Chen (2023), we investigate the effects of decreased shipping costs resulting from the consolidation of transport networks on downstream firms. However, transport companies and their downstream customers typically establish long-term relationships, and our knowledge of these relationships is limited. To shed light on these issues, we examine how these relationships are formed and how demand and supply shocks affect the formation of such relationships. One challenge in examining these relationships is that, typically, the specifics of such contracts are not observed. However, the U.S. natural gas pipeline industry is an exception, as all contract details are observable. This includes information on the ownership of transmission capacity in the primary market and details of transactions related to capacity release in the secondary market. By examining the formation of long-term relationships in the natural gas pipeline industry, we offer insights that could be applied to other similar transport markets with substantial sunk investment costs.

Our study of the secondary market in the natural gas pipeline industry also contributes to a literature investigating the effect of capacity constraints, bottleneck, and congestion on market efficiency measures such as price integration across hubs (Avalos, Fitzgerald and Rucker, 2016, Brown and Yücel, 2008, Marmer, Shapiro and MacAvoy, 2007, Oliver, Mason and Finnoff, 2014). This literature studies physical constraints and how regulation interacts with these physical constraints in preventing spatial price integration. By contrast, we study the role of artificial capacity constraints created by long-term contractual agreements that reserve capacity, which has not been studied before. Our findings on the effectiveness of the secondary market suggest that these constraints are important and should be considered along with physical constraints when studying market efficiency.

Lastly, our findings on the exercise of market power in the secondary markets relate to the literature on intermediation in search markets and monopsony power (Berger, Herkenhoff, Kostøl and Mongey, 2023, Gehrig, 1993, Salz, 2022, Spulber, 1996). Our results are consistent with some of these results in contexts where intermediaries (marketers) have market power.

4.2 Industry Background

Natural gas transmission pipelines provide for the bulk of natural gas transportation within North America, often transporting commodities hundreds or thousands of kilometers from production locations to three main demand markets: electricity generation, local distribution companies (LDCs) primarily used for space heating and industrial consumption. In comparison to other modes of domestic freight, pipelines transport the most tonnes-kilometers. Natural gas transportation in anything other than pipelines is complicated and often unprofitable.

Within the U.S., there are more than 120 major natural gas transmission pipelines whose tolls/rates are directly regulated by the Federal Energy Regulatory Commission (FERC). Most of these pipelines facilitate the transmission of natural gas from production areas, concentrated in regions such as Texas, Pennsylvania, and Canada, to demand areas such as major metropolitan areas, major export interconnections, and natural gas storage facilities. The major interstate transmission pipelines are shown in the map below in Figure 4.1.

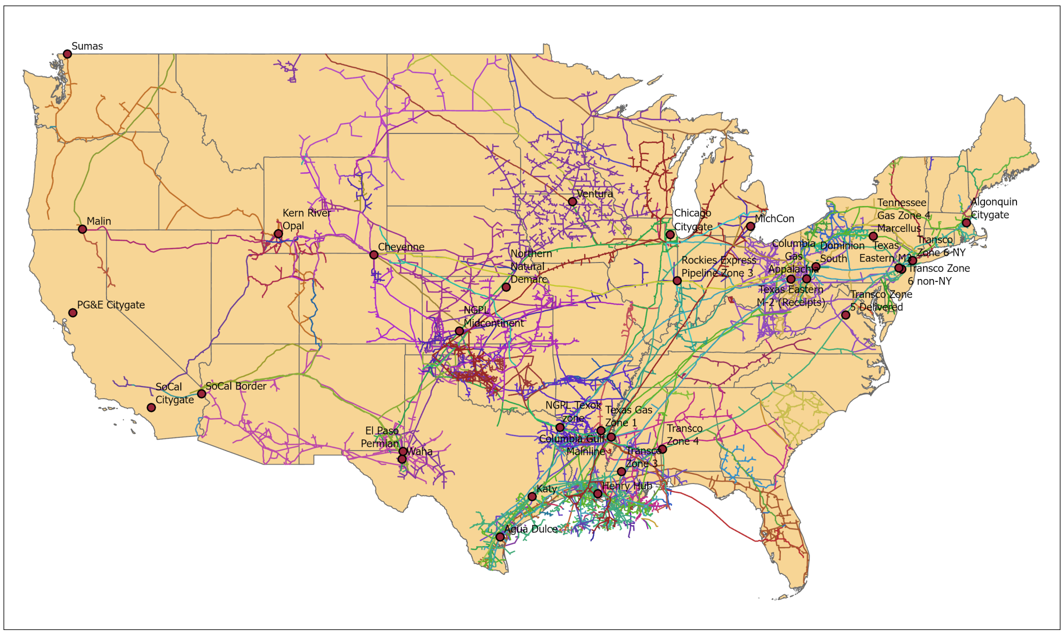

Figure 4.1: Natural Gas Pipelines in the U.S.

Transmission pipelines are extremely costly to build, averaging $5 million/km with 10 to 20 years for the pipeline to break even on its investment. To share the risk associated with large sunk investment costs, the pipelines are required to sign long-term contracts (typically ten years or more) with customers prior to the pipeline construction, which entitles the customer to reserved capacity. The requirements for long-term contracts come from both sides of the market. Pipeline companies are required to sign these contracts to obtain funding. Large customers, such as power plants and local distribution companies (LDCs), are required by regulatory agencies to have enough available capacity to sustain through unexpected demand shocks. Smaller natural gas users do not contract capacity with pipelines directly and typically get gas from marketers.

Long-term contracts in the natural gas transmission pipeline industry primarily focus on shipment capacity,[2] as shipment pricing is often heavily regulated. Pipeline rates are often set to enable pipelines to recover all prudently incurred expenses associated with service delivery while also earning a reasonable profit. Regulators usually establish this by determining the revenue requirement for pipelines — the yearly revenue needed to maintain service and secure a fair return.[3] In Appendix I, we provide further details on rate regulation within the pipeline industry.

While long-term contracts in the primary market help reduce investment risks, they can introduce substantial frictions in capacity allocation amidst fluctuations in the natural gas supply and demand, thereby obstructing gas delivery to customers in need. The secondary market has emerged to address these frictions. In this market, service requesters who hold long-term contracts in the primary market can trade their contractual rights. This process, known as "Capacity Release," allows for the sale of all or part of a contract holder's rights for varying durations, ranging from less than a month up to the entire contract length.

Figure 4.2 illustrates the interactions between the primary and secondary markets. Within the capacity release market, those holding primary contracts are known as *releasers*, and their counterparts are referred to as *replacers*.

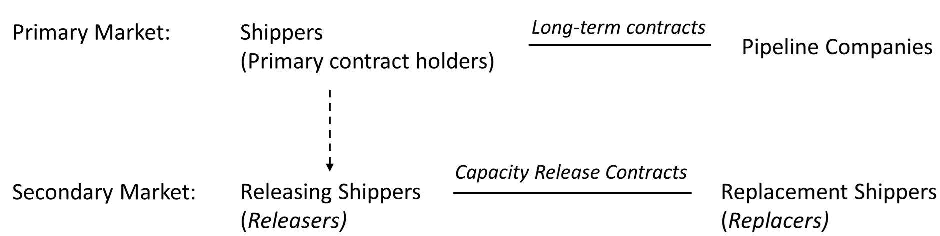

Figure 4.2: Relationships of the Primary and the Secondary Market

[2] A firm transportation contract grants capacity to a service requester at one or more points along a pipeline. Capacity is either specific as to both location (point) and quantity or is general as to location and specific as to quantity. A firm transportation contract gives a service requester the right to cause a TSP to receive a specific quantity of gas from that service requester at a point and/or deliver a specific quantity of gas to that service requester at a point over a specific time period.

[3] A typical revenue requirement will be comprised of the following:

$$Operations\&Maintenance\ Costs + Taxes + Depreciation + Return = Revenue\ Requirement$$

Looking more specifically at the "fair" return for a pipeline in the cost of service regulation terminology, the return is typically determined through a rate base. A rate base is defined as a pipeline's gross plant in service less its accumulated depreciation. You then earn a return based on an assumed capital structure the regulator sets. Those portions of debt and equity then earn a return based on your cost of debt and cost of equity, both of which are determined by the regulator. The cost of debt is typically based on a pipeline's outstanding debt. In contrast, the cost of capital is determined through regulatory proceedings utilizing various financial models to determine an appropriate cost of equity for the pipeline based on similarly risky assets (mainly other pipeline systems).

For primary contract holders, the value of transportation tends to rise when there are substantial price differentials between natural gas trading hubs. Consequently, we expect to observe increased capacity releases from these customers when such differentials widen. Additionally, it is important to note that many natural gas end-users either do not possess transportation capacity on interstate pipeline systems or are not connected to natural gas distribution utilities. These customers are often industrial facilities, agricultural operations, smaller power generation facilities, and natural gas retailers. Thus, many of these natural gas end users rely on natural gas marketers to meet their gas needs in terms of supply and transportation.

These end users lack pipeline transportation capacity mainly due to concerns about creditworthiness and balance sheet obligations. The credit obligations necessary for a company to hold pipeline transportation services are often steep, and most companies cannot meet them. For example, a typical credit evaluation criterion for firm service on a natural gas pipeline is to provide security guarantees for three months of firm service at the maximum tariff rate for the entire volume of your contract. This requires companies to have large amounts of cash on hand (in the form of an advance deposit), a strong standing letter of credit from a financial institution, an acceptable security interest in collateral, or a guarantee from a more credit-worthy parent company. As for the balance sheet obligations, given the take-or-pay nature of natural gas transportation firm service contracts, financial institutions view these transportation contracts as debt obligations. If a company were to take out large amounts of transportation capacity, this would result in a large liability on their balance sheet, which may impact their credit metrics, impacting their ability to secure their own financing and financial obligations.

Given these restrictive requirements, many of these end users rely on marketers to arrange for the supply and transportation of their gas needs. This comes at an increased cost to the end user, as marketers often require a service fee or markup for arranging the supply and transportation of natural gas.

Since these end users rely on marketers to provide them service when unexpected shocks in demand happen, either to end-use residential, commercial, and industrial demand (increased demand for natural gas retailers) or unexpected shocks in their various industries that do not impact demand for gas utilities/retailers (for example an increase in demand for steel production) they often turn to marketers to supply them with additional natural gas. While the primary market for natural gas transportation is mainly held by utilities to meet their regulatory obligations, they do not require the full use of their transportation capacity for most periods of the year. They would prefer to release that capacity to a marketer who will ultimately provide the transportation services to another end user.

4.3 Data

We utilize four datasets in our analysis—first, the index of customer data. An index of customers provides each specific contract on a pipeline system, including which shipper holds that contract, the start and end dates of the contract, the type and path of the service, the amount of pipeline capacity the shipper can utilize, and the contract rate. In the United States, pipelines regulated by the Federal Energy Regulatory Commission (FERC) are required to publish an index of customers from which we obtain our data. Second, the capacity release data from each pipeline system, which includes details such as the transaction date, type of contract, duration of the contract, rate, and options associated with the contract (such as recall or reput information).[4] Third, the daily spot prices of natural gas at all US-based hubs, which Capital IQ Pro compiles. We obtained the first three datasets from 2004 to 2023.[5] Last, we use geographic data from the U.S. Energy Information Administration (EIA) to gather information about the entire U.S. interstate natural gas pipeline system and the locations of major natural gas gateways and hubs.

4.3.1 Descriptives for the Primary Market

We first provide summary statistics for the long-term contracts in the primary market for selected pipelines. Across our sample, the average duration of a contract is approximately nine years.

Table 4.1: Contract Duration (In Months)

Pipeline	p5	p25	p50	p75	Mean	Max
El Paso Natural	9	57	115	180	122	475
Natural Gas Pipe	12	36	60	120	84	432
Texas Eastern	13	48	105	184	118	612
Transwestern	4	12	36	101	65	361
All Sample	12	36	72	154	100	612

Next, we provide insight into the competitive landscape of the industry. Table 4.2 shows the overview of market concentration in U.S. Interstate Pipelines.

[4]With the advent of the capacity release market, the FERC required pipelines to openly post the deals that their service requesters were seeking to transact (FERC Order No. 636, et al.).

[5]While all of this data is technically public information, most pipelines periodically delete information from past transactions, such as contracts signed in the capacity release market. We purchase the entire history of the capacity release market and spot prices at hubs between 2004 and 2023 from the Capital IQ Pro database to access historical data.

Table 4.2: U.S. Market Share: Top 5 Interstate Pipelines by Capacity (2022)

Pipeline	**Market Share**
Transcontinental Gas P L Co	.1
Texas Eastern Trans Corp	.08
Tennessee Gas Pipeline Co	.08
ANR Pipeline Co	.05
Rockies Express Pipeline	.04
Mean	.006
N	170

Table 4.3 shows the major contract holders in the primary market. The data indicate that the top 5 companies own 75% of interstate pipeline capacity. The market power observed in the primary market will significantly impact who can participate in the capacity release market (the secondary market) and will affect shippers' decisions to release capacity during demand or supply shocks.

Table 4.3: Top 10 Contract Holders by Capacity in the Primary Market (2022)

Primary Owner	**Market Share**
Kinder Morgan	.195
TC Energy	.192
Enbridge	.151
The Williams Companies	.12
Energy Transfer	.1
Tallgrass Energy	.05
Boardwalk Pipeline	.035
Berkshire Hathaway	.032
Boardwalk Pipelines	.024
Dominion Energy	.024
Mean	.013
N	76

4.3.2 Descriptives for the Secondary Market

Regarding the capacity release market, we first demonstrate its importance. Table 4.4 shows the amount released in the secondary market compared to the primary market. The percentage of the amount released is calculated at the Pipeline-Shipper-Quarterly level, using the Max Daily Transport quantity as a basis. The result shows that in any quarter, on average 30% of capacity is being treaded on the secondary market.

Table 4.4: Amount Released in Comparison to IOC Data (Quarterly)

	p25	p50	p75	Mean	Max
Proportion of quantity released	2.3%	6.8%	25.0%	32.4%	4212.7%

Next, we examine the percentage of shippers participating in the capacity release market at any point during the contract period. Table 4.5 reveals that 40% of shippers holding long-term contracts participate in the capacity release market. This underscores the significant importance and active usage of the secondary market in this industry. The data therein highlights substantial variations in both the number of capacity releases and the number of participants in the secondary market across different pipeline systems every quarter.

Table 4.5: Number of Shippers That Are Within the Secondary Market

Not in the Capacity Release Market	555	60%
In the market	373	40%
Total	928	100%

Table 4.6 shows the average duration of a contract in the capacity release market. Most contracts last a month, with some lasting half a year or a whole year. It is rare for a contract in the secondary market to last for more than a year. Figure 4.3 illustrates the distribution of contract durations.

Table 4.6: Contract Durations in the Capacity Release Market (Days)

	Mean	p5	p50	p95	Min	Max
Contract Duration	85.17	27	30	364	1	1000

In the capacity release market, there are several options that a contract might include: recallable, reputable, capacity resale allowed, affiliation of the counterparty with the contracting party, and inclusion of previously released capacity. Table 4.7 shows the percentage of contracts that feature these options.

We then present data descriptives to illustrate the interactions between releasers and replacers in the secondary market. The selected pipelines are “El Paso Texas Pipeline”, “Texas Eastern Transmission”, “Natural Gas Pipeline Company of America” (*NGPL*), and the Transwestern Pipeline. We observe data between 2006 and 2023. Using the points identifiers for the contracts, we were able to separate capacity release between different categories:

Table 4.7: Percentage of Contracts With Different Options

Year	Recallable	Reputable	Resale Allowed	Affiliate	Previously Released
2006	72%	58%	72%	0%	13%
2011	93%	65%	92%	1%	12%
2016	97%	78%	97%	0%	6%
2021	98%	78%	96%	0%	18%

Table 4.8: Distribution of sub-contracts by category—Selected pipelines

	Number of sub-contracts	Percentage
Compressor	984	0.7
Delivery to End User	3,057	2.1
Delivery to an LDC	**58,877**	**40.3**
Exchange	59	0.0
Gas Processing Plant	458	0.3
Gathering	278	0.2
Interconnect	**53,14**6	**36.4**
LNG	36	0.0
Park and Loan	1	0.0
Pool	3,662	2.5
Power Plant	93	0.1
Receipt by LDC	79	0.1
Segment	313	0.2
Stand Alone Meter	3,033	2.1
Storage Injection	264	0.2
Storage Quantity	**19,093**	**13.1**
Storage Withdrawal	352	0.2
Unknown	973	0.7
Wellhead	1,265	0.9
Total	146,023	100.0

Notes: A sub-contract is defined as an agreement on capacity release at a specific point, as part of larger capacity release contracts.

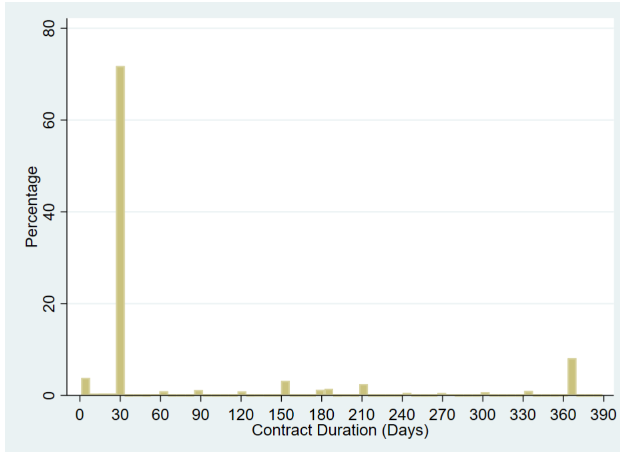

Figure 4.3: Distribution of Contract Durations

We aggregate these categories into local distribution companies (LDC), storage, and others (where *others* are mostly comprised of transmission between pipelines). Additionally, we separated contracts into two types based on the frequency of contracts between the same shippers on a pipeline: (1) Met < 5 times and (2) Met ≥ 5 times.

We refer to shippers that met more than five times as engaged in long-term relationships. These categories will help us understand the results of the event studies.

4.4 Emprical Model and Results

4.4.1 Event Study — Investigation of a Cold Wave

February 2021 Cold Wave

In February 2021, an extraordinary cold wave, driven by a polar vortex's southward shift following a sudden stratospheric warming, swept across Canada, the United States, and northern Mexico. This rare meteorological event led to severe winter storms, unprecedented snow, and cold temperatures in states forming the deep south, such as Texas, Oklahoma, and Arkansas.

This cold wave led to an unexpected demand increase for home heating, affecting demand for natural gas as a direct heat source for homes and indirectly through electricity generation. In states such as Texas, natural gas directly heats 35% of homes, and electricity heats the remaining 65%. Meanwhile, natural gas is responsible for over 50% of electricity generation.[6] We expect this cold wave to have led to unprecedented pressure on natural gas pipelines, not only due to

[6]Texas uses natural gas for electricity generation and home heating.

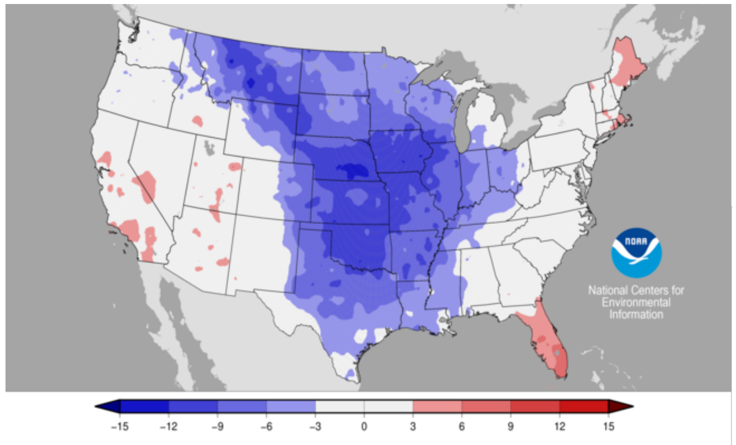

Figure 4.4: Temperature Deviation From Historical Average – February 2021

Notes: Temperatures are expressed as deviation (in Fahrenheit) from average temperatures during the 20^{th} century across 5km grid points. Source: National Center for Environmental Information (NCEI).

the unexpected rise in demand but also as some pipeline segments froze and burst, disrupting some of the supply.[7]

In this section, we investigate the secondary market's role in allocating natural gas capacity during the cold wave, some of the mechanisms underlying this effect, and potential trade-offs associated with the secondary market. We first investigate a simple difference-in-difference event study that allows us to compare outcomes between treated and control pipelines throughout the cold wave. Outcomes such as contracted quantity may have persisted in the weeks following the shock because the end date was not known ex-ante.

Defining Control and Treatment Pipelines

To define control and treatment pipelines, we looked at interstate pipelines with a similar activity ratio in the secondary market relative to the pipeline's total capacity and location. We chose El Paso Natural Gas Company as a treated pipeline, which transports gas from the San Juan, Permian, and Anadarko basins to selected states, including Texas and Oklahoma. The El Paso pipeline is a good candidate because it covers most hubs that saw massive hikes in natural gas prices relative to prices at Henry Hub during the cold wave. Prices in southwestern Texas and southeastern California rose between 10 and 25 times the national benchmark due to the unprecedented uptake in demand. Panel (b) of Figure 4.5 shows the map of the treated pipeline with selected hubs.

As a control, we chose the Gas Transmission Northwest pipeline (GTN), which takes gas from

[7]February 2021 North American cold wave.

(a) Control Pipeline (b) Treated Pipeline

Figure 4.5: Treated and Control Pipeline Maps With Selected Hubs

Western Canada and takes it to Northern California—in contrast to abnormal weather variation in Texas, Figure 4.4 indicates that Northern California was not affected by the cold wave because temperatures did not deviate from historical averages. Panel (a) of Figure 4.5 shows the control pipeline map.

To further validate the categorization of treated and control groups, we calculate the price of natural gas at selected hubs served by treated and control pipelines relative to the Henry Hub benchmark. Figure 4.6 shows the prices. We can see that during and after the cold wave, the relative prices at control hubs are quite stable. However, for treated hubs, there is a spike in gas prices during the cold wave.

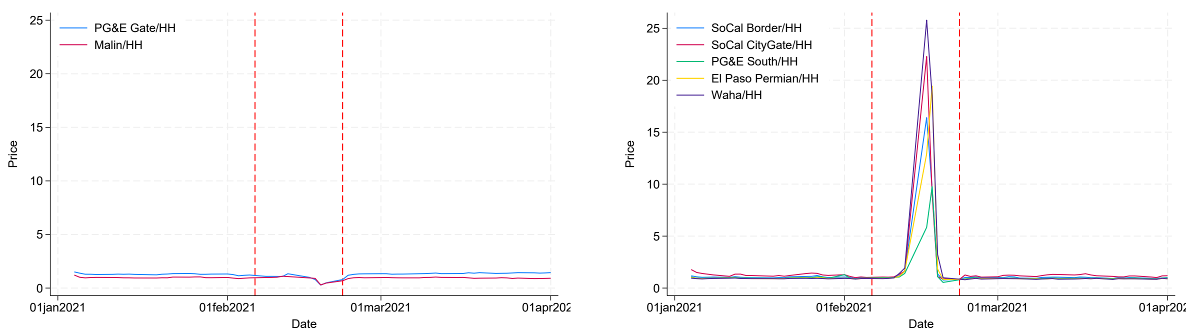

(a) Relative Price at Control Hubs (b) Relative Price at Treated Hubs

Figure 4.6: Prices Relative to Henry Hub

Notes: This figure plots the price of natural gas at selected hubs served by treated pipelines relative to the Henry Hub benchmark. The dotted red lines correspond to the official beginning of the cold wave (February 6^{th}, 2021) and the end of the cold wave (February 22^{nd}, 2021).

As discussed by Marmer, Shapiro and MacAvoy (2007), the possibility for arbitrage could imply that the entire pipeline network is treated as intermediaries buy gas from unaffected regions (here, the West Coast) and sell it in affected regions (the Deep South). However, such arbitrage would increase demand at unaffected hubs and raise gas prices. We do not see a large price rise at unaffected hubs such as Malin and PG&E Citygate (which serves San Francisco). Just as Marmer et al. (2007) found, there are likely important bottlenecks in the system preventing arbitrage between regions far away from each other. While the El Paso pipeline is a larger pipeline system than the GTN pipeline, the average percentage of capacity from the capacity release market is similar across both pipelines.

4.4.2 Difference-In-Difference Analysis

We perform a difference-in-difference event study with a single treated pipeline and a single control pipeline, allowing us to investigate the effect using a parsimonious two-way fixed-effects specification. We cannot rule out anticipation as shippers try to predict adverse weather events. Such extreme events may be correlated with temperature patterns over an extended period, such as an abnormally cold winter. We also cannot rule out persistent effects because shippers make decisions on capacity release without knowing the end date of the demand shock and may want to air on the side of caution in the aftermath of the cold wave. Polar vortices can last from a few days to multiple months.

We allow for anticipation and persistence by investigating the effect in a larger window around the cold wave. As Figure 4.7 suggests, capacity released in the secondary market during the cold wave significantly increased in the El Paso pipeline (by 30% immediately following the beginning of the cold wave). It did not increase in the GTN pipeline. However, there was also a significant increase in capacity release for El Paso relative to GTN one month prior to the cold wave, which persisted up to one month after the cold wave.

We also need to account for important seasonal variations across both pipelines. A careful look at the time series of capacity released in the secondary market outside of the window around the cold wave suggests a very distinct time series across pipelines. Much of the variation around the cold wave may obfuscate important seasonal patterns. Much of this variation can be attributed to pipeline-specific seasonal variation, in which different shippers make recurring contracts with each other. For example, utilities may lease some of their capacity to power plants every summer.

As shown in the main specification below, we allow for heterogeneous seasonal variation across pipelines to account for this rich heterogeneity. To study the causal effect of the 2021 cold wave on capacity released, we then propose the following two-way fixed effects specification:

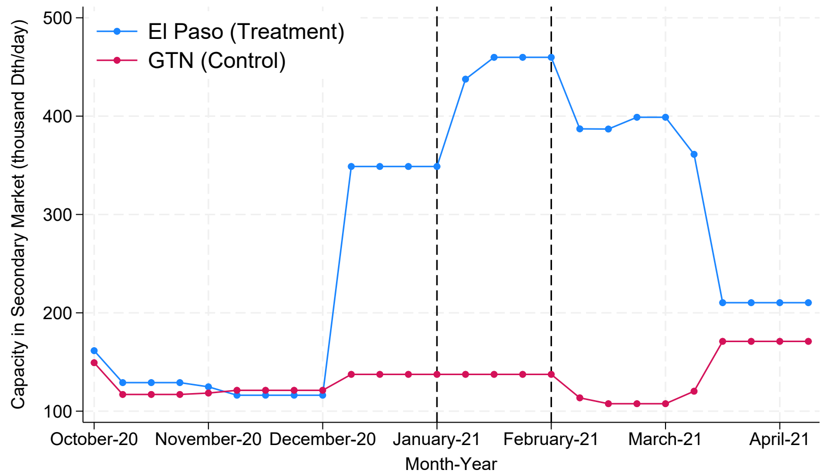

Figure 4.7: Total Capacity in Secondary Market by Pipeline

Notes: The outcome variation displayed in this figure is the total capacity available in the secondary market each week. It is constructed by aggregating all contracts signed on previous dates with contracted capacity released during the current week. In the time axis, each month's tick corresponds to the end of the month. The dashed black lines correspond to the cold wave's beginning and end.

$$y_{it} = \alpha_i + \alpha_t + q_{it}^T\delta + \sum_{m=-5}^{10} \beta_m D_{i,t-m} + \epsilon_{it} \tag{4.1}$$

Where i indexes pipeline, t indexes weeks. y_{it} is the log of capacity available in week t from the secondary market. q_{it} are pipeline-specific seasonality controls. These include week of the year (1-52) and year of observation. $D_{i,t-m}$ is a treatment indicator defined as follows:

$$D_{i,t-m} = \begin{cases} 1 & \text{if } i \text{ is treated and we are } t-m \text{ weeks relative to start of cold wave} \\ 0 & \text{otherwise} \end{cases}$$

The outcome variable y_{it} aggregates all contracts signed at previous dates $t-k$ that have contracted capacity to be released at t. Since most contracts are signed for one month, there is mechanically a lot of auto-correlation in y_{it}. We cluster standard errors at the monthly level to conduct autocorrelation-robust inference within 30-day periods.

To visualize the results, we follow the approach of Freyaldenhoven, Hansen and Shapiro (2019) by normalizing the treatment effect one week before the shock, plotting the cumulative treatment effect in a window around treatment. Under a null hypothesis that the cold wave was not anticipated, lasted throughout February (4 weeks), and did not have persistent effects, the cumulative treatment effect would be defined as:

$$\gamma_k = \begin{cases} 0 & \text{for } k < 0 \\ \sum_{m=0}^{k} \beta_m & \text{for } 1 \leq k \leq 4 \\ 0 & \text{for } k > 4 \end{cases}$$

However, as discussed earlier, shippers may anticipate an abnormally cold winter as they predict natural gas requirements before the winter starts. As discussed previously, there is likely much persistence in these treatment effects since contracts are signed without knowing when the cold wave will end. For these reasons, we plot an estimate of the cumulative effect up to five weeks before and ten weeks after the beginning of the cold wave, along with 95% confidence intervals. We then define estimates of the cumulative event path as follows:

$$\hat{\gamma}_k = \begin{cases} \sum_{m=-k}^{-1} \hat{\beta}_m & \text{for } -5 \leq k < 0 \\ \sum_{m=0}^{k} \hat{\beta}_m & \text{for } 0 \leq k \leq 10 \end{cases} \tag{4.2}$$

To investigate anticipation and persistence more formally, we also present the results from a specification that aggregates the effect up to 2 months prior to the cold wave and two months after the cold wave using the following specification:

$$\begin{aligned} y_{it} = \alpha_i + \alpha_t + q_{it}^T \delta + \underbrace{\beta_d D_{i,t \in dec20} + \beta_j D_{i,t \in jan21}}_{\text{Treated pipeline two months before}} + \beta_f D_{i,t \in feb21} \\ + \underbrace{\beta_m D_{i,t \in mar21} + \beta_a D_{i,t \in apr21}}_{\text{Treated pipeline two months after}} + \epsilon_{it} \end{aligned} \tag{4.3}$$

4.4.3 Results — El Paso (Treatment) and GTN (Control)

Figure 4.8 shows the estimated cumulative event path from the main specification in Equation 4.1. While we do not find statistically significant violations of the pre-trend assumption, these results do not suggest a lack of anticipation. Rather, it should be interpreted as a lack of anticipation that the cold wave would start earlier than planned. Anticipation is directly embedded in the outcome variables, aggregating previously signed contracts. On the other hand, we find large and significant persistent effects many weeks after the cold wave, suggesting that shippers did not know when the demand shock would end. Overall, the cold wave caused an increase in capacity release by an average of 50% across the duration of the cold wave. See Table 4.9.

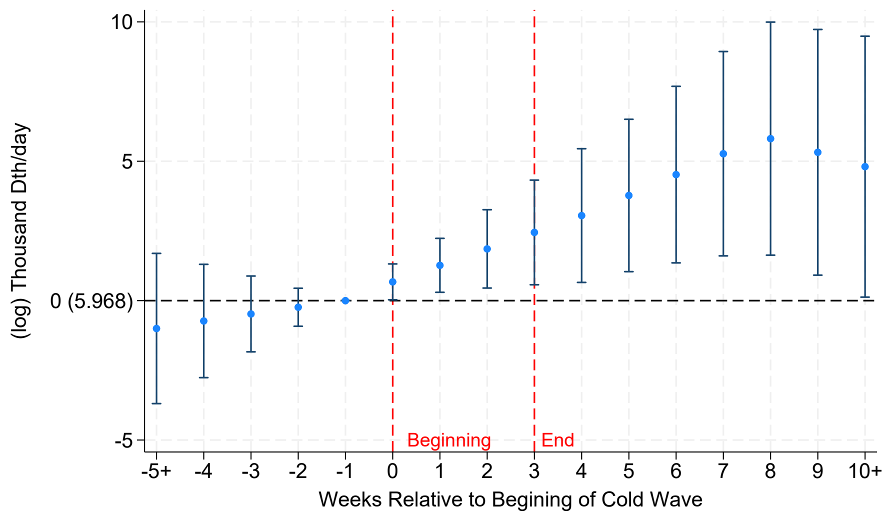

Figure 4.8: Baseline Cumulative Event Path

Notes: This figure plots the event path specified in equation 4.2 along with 95% confidence intervals. Specifically, we plot the cumulative effect of the February 2021 cold wave on the quantity of pipeline capacity contracted in the secondary market. Specifically, this quantity is defined as the total quantity released in the current week, aggregating all contracts signed prior to and up to the current week. We do so because many contracts are often signed weeks before the capacity is released. We exclude contracts that are signed for periods of one year or more.

To further validate these findings, we performed a robustness check by imposing a placebo cold wave one year before, in 2020, and found no evidence of a placebo treatment effect.

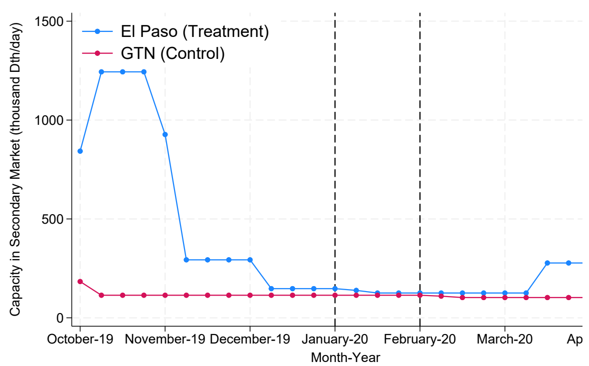

(a) Total Capacity In Secondary Market by Pipeline

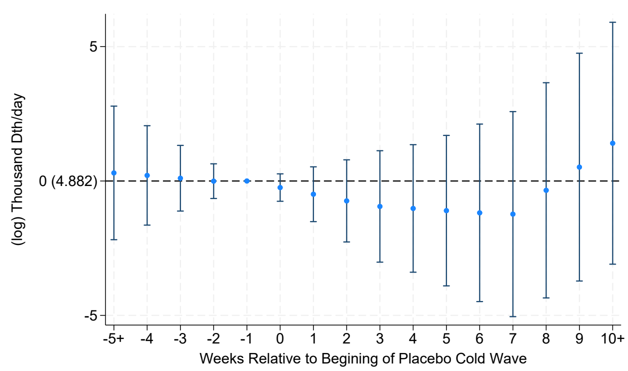

(b) Cumulative Event Path

Figure 4.9: Placebo Cold Wave (One Year Before)

Secondary Market and Search Costs

To investigate trade-offs associated with a deregulated secondary market, we first provide evidence that the cold wave, interpreted as a demand shock, increased the search cost between

Table 4.9: Main Results — Average Effect of Cold Wave on Capacity Release

	Real cold wave (log) Capacity	Placebo cold wave (log) Capacity
$\mathbb{I}$(Two Months Before x Treated)	0.251	0.349
	(0.376)	(0.405)
$\mathbb{I}$(One Month Before x Treated)	0.215	-0.112
	(0.336)	(0.331)
$\mathbb{I}$(During x Treated)	0.579**	-0.203
	(0.237)	(0.242)
$\mathbb{I}$(One Month After x Treated)	0.674**	-0.130
	(0.280)	(0.278)
$\mathbb{I}$(Two Months After x Treated)	-0.266	0.624**
	(0.338)	(0.294)
N	1,278	1,278
adj. R^2	0.784	0.782

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: these results are exactly as specified in Equation 4.3. As such, they are not the cumulative treatment effects but the average treatment effect across the specified period. The first column presents results of the real cold wave, whereas the second column presents results for the placebo cold wave (one year earlier), including two months before and after this placebo cold wave.

a releasing and a replacement shipper. As argued previously, the secondary market provides a mechanism to reduce these search costs by fostering the formation of long-term relationships. Indeed, more than 95% of contracts in the secondary market are private arrangements between shippers, and shippers often use these private arrangements to form long-term relationships. We now argue that these long-term relationships are crucial when shippers face large unexpected demand shocks such as the cold wave of 2021.

To be consistent with our classification in Section 4.3, we group all contracts from the secondary market into two categories: contracts between shippers that met less than five times in the past (which we call new relationships) and contracts between shippers that met more than five times in the past (which we call long-term relationships). Below is the distribution of relationship frequency between shippers:

We then re-estimate the two-way fixed effects regression of equation 4.1 separately by groups. When grouping the contracts by relationship frequency in Figure 4.10, we find that the effect of the cold wave on capacity released is overwhelmingly driven by long-term partners. Moreover, the effect nearly doubles when narrowing down on long-term partners. On average, we find that the cold wave caused an increase in capacity between long-term partners by approximately

Table 4.10: Distribution of Relationship Frequency by Pipeline

Frequency of Identical Shippers Signing Contracts	El Paso (Treatment) *Percentage (%) of Contracts*	GTN (Control) *Percentage (%) of Contracts*
Never met	1	2
Met > 0 but ≤ 5 times	22	9
Met > 5 but ≤ 10 times	18	16
Met > 10 but ≤ 15 times	15	10
Met > 15 but ≤ 20 times	8	5
Met > 20 but ≤ 25 times	9	25
Met > 25 but ≤ 30 times	7	5
Met > 30 times	20	28
Total	100	100

Notes: The distribution corresponds to the percentage of unique contracts signed between shippers.

90%.

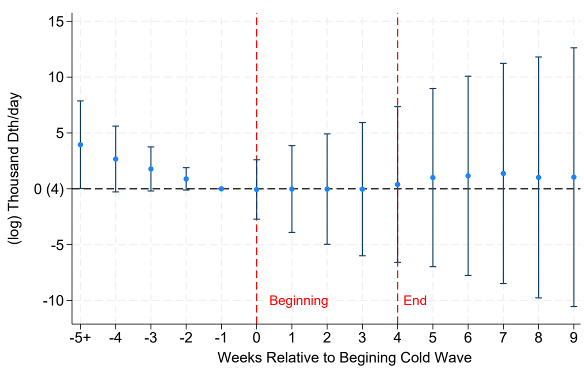

(a) Shippers Met < 5 Times

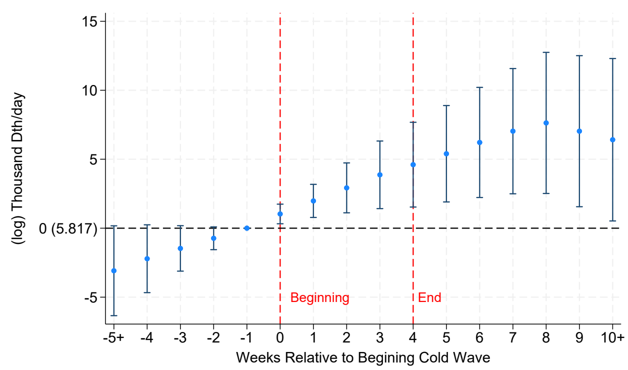

(b) Shippers Met ≥ 5 Times

Figure 4.10: Cumulative Event Path by Relationship Length

Finding new partners has a high search cost, and the opportunity cost of searching is high in times of unexpected shock. For this reason, shippers reduce the search cost by leveraging their existing relationships. Moreover, the secondary market provides a platform for shippers to form new relationships during normal times, which can be leveraged during shocks. In this context, the secondary market provides a mechanism to reduce the search cost. Due to the FERC regulation, which mandates that all contracts be posted on public platforms, the secondary market fosters the formation of long-term relationships by providing shippers with daily information on the key players in each market.

We further substantiate this narrative that the cold wave is associated with a higher opportunity cost of searching by looking at the characteristics of contracts during the cold wave. Indeed, what drives the main result is not an increase in the number of contracts, which would be costly, but rather an increase in the average quantity released by contract.

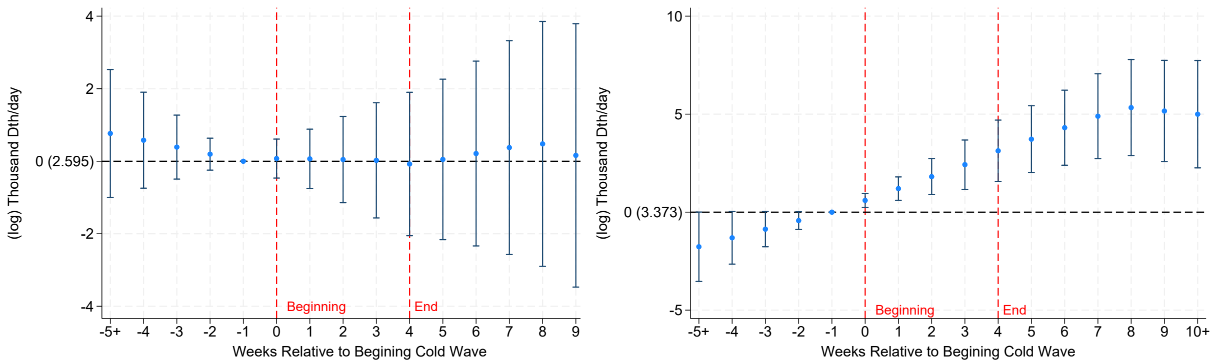

(a) Number of Contracts (b) Quantity per Contract

Figure 4.11: Event Path — Number of Contracts and Quantity per Contract

Table 4.11: Average Effect of Cold Wave — Extensive and Intensive Margins

	(log) Number of Contracts	(log) Average Quantity per Contract
$\mathbb{I}$(Two Months Before x Treated)	1.146***	-0.895***
	(0.141)	(0.290)
$\mathbb{I}$(One Month Before x Treated)	-0.194	0.409*
	(0.125)	(0.229)
$\mathbb{I}$(During x Treated)	0.003	0.577***
	(0.126)	(0.166)
$\mathbb{I}$(One Month After x Treated)	0.084	0.589***
	(0.152)	(0.174)
$\mathbb{I}$(Two Months After x Treated)	-0.229	-0.037
	(0.160)	(0.210)
N	1,278	1,278
adj. R^2	0.726	0.860

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

While this narrative on search costs is broadly consistent with the literature, there is one key difference. In the literature on intermediation with search costs (Gehrig, 1993, Spulber, 1996), it is often assumed that parties on two sides of a market can either engage in search or go through an intermediary to avoid search costs. In our case, the margin to reduce search cost is forming relationships between an intermediary (the replacer) and a releaser. This means that the presence of an intermediary does not eliminate search costs. We also argue that both releasers and replacers engage in search. The replacers search for available capacity when they face arbitrage opportunities between hubs. Releasers search for buyers when they do not need capacity to avoid paying reservation fees for unused capacity.

Market Power

While the secondary market provides a platform for shippers to reduce search costs of releasing capacity, it is also largely unregulated. For the majority of contracts signed, releasing and replacement shippers can freely negotiate a contract rate, which can be below, equal, or above what the releaser pays to the pipeline.[8] Deregulation can incentivize market participants to engage in the secondary market but can also introduce market power, causing friction. We now investigate market concentration across different pipelines and zones. We find evidence that some markets have more concentrated releasers, while some have more concentrated replacers, and some have both. We relate this variation in market concentration to capacity release prices and find that releasers (replacers) with a higher market share tend to charge higher (lower) prices.

Table 4.12: Summary Statistics — Market Power and Prices (Secondary Market)

	Replacer HHI	Releaser HHI	$P < Max$	$Price = Max$	$P > Max$	$P = 0$
El Paso	1,660	4,821	0.81	0.15	0.04	0.21
GTN	1,723	4,746	0.21	0.75	0.04	0.03
GL	5,366	5,491	0.48	0.51	0.01	0.33
NGPL	1,529	1,769	0.68	0.24	0.08	0.09
Texas Eastern	2,708	1,586	0.50	0.42	0.08	0.32
Transwestern	1,889	2,685	0.95	0.03	0.02	0.15

Notes: In this Table, P refers to the tariff charged for a specific section of a contract in the secondary market, and always has two parts: a reservation fee and a volumetric charge based on gas flown. For example, a shipper can release capacity between delivery point A and receipt point B and between delivery point B and receipt point C. In this case, there would be four prices. Max corresponds to the maximum tariff set by the FERC for those four prices. HHI refers to the Herfindahl–Hirschman index.

Table 4.12 summarizes market concentration across selected pipelines with considerable activity on the secondary market. We find substantial variation in market power across pipelines. For example, the Great Lakes (GL) pipeline has a high concentration of replacers and releasers, whereas El Paso has a low concentration of replacers but a high concentration of releasers. We also find lots of variation in average tariffs relative to the tariffs. Since the FERC sets these maximum tariffs specific to delivery and receipt points, they vary widely across contracts and provide a benchmark for comparing tariffs. However, Table 4.12 does not suggest a monotonicity between releaser/replacer market power and prices. For example, both El Paso and GTN

[8] There are some exceptions. Certain contracts cannot have a rate above some maximum threshold the FERC sets. This constraint is based on the contract duration. Tariffs on capacity released for 31 days or less or one year or more can vary freely. Tariffs on capacity released between 31 days and one year cannot exceed the maximum tariffs. As Figure 4.3 suggests, the vast majority of contracts fall within the fully deregulated time constraints.

have a high concentration of releasers relative to replacers, and tariffs are significantly more likely to be below the maximum in El Paso.

To understand this point better, we also show in Figure 4.12 that the cumulative distribution of market power in releasing shippers tends to be highly correlated with the cumulative distribution of market power for replacement shippers, which is likely to be correlated with market thickness. Great Lakes is the most concentrated in both releasers and replacers, whereas NGPL is the least concentrated. GTN and Transwestern are more concentrated in releasers and replacers than El Paso and Texas Eastern. For this reason, only looking at variation in market concentration may not provide enough variation to investigate its effect on tariffs.

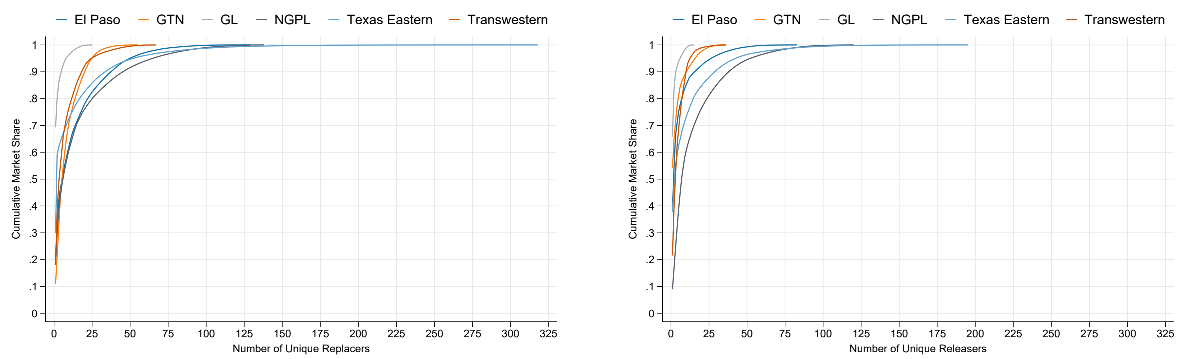

(a) Replacers (b) Releasers

Figure 4.12: Market Concentration Across Pipelines

Next, we investigate the role of releaser and replacer market power on tariffs in more detail. Specifically, we investigate variation in market power at the individual releaser/replacer level. We specified the following regression model, where we consider both the log of tariffs and the log of tariffs relative to maximum tariffs as outcome variables:

$$\ln Y_{pcit} = \alpha_t + \alpha_p + X_{it}^T\beta + z_{cpit}^T\gamma + \beta^{rep}\ln S_{rep,it} + \beta^{rel}\ln S_{rel,it} + \epsilon_{cpit}$$

Where p indexes a delivery and receipt point pair, c indexes a contract, i indexes a pipeline, and t indexes days. The main independent variables of interest are the market share of releasing and replacement shippers signing the contract, $\ln S_{rep,it}$ and $\ln S_{rel,it}$ respectively. For robustness, I consider various specifications for this market share: (1) total pipeline market share (constant across all years); (2) annual pipeline market share; (3) total market share in pipeline-specific zones (constant across all years); and (4) annual market share in pipeline-specific zones.

X_{it} are seasonal adjustments that vary by pipelines, and Z_{it} are control variables and include relationship length between shippers and other contract characteristics such as contract duration. α_p are pipeline-specific delivery-receipt point fixed effects, so they absorb pipeline fixed effects, and α_t are week fixed effects. Here, β^{rep} and β^{rel} capture variation in market share across shippers and within shippers over time because we omit shipper fixed effects.[9] All the results are robust to the inclusion/exclusion of all control variables, fixed effects, and seasonality.

Table 4.13: Estimation Results — Log Tarrifs

	(1)	(2)	(3)	(4)
Releaser share	0.088***			
	(0.003)			
Replacer share	-0.040***			
	(0.001)			
Annual releaser share		0.063***		
		(0.003)		
Annual replacer share		-0.042***		
		(0.001)		
Zone releaser share			0.113***	
			(0.003)	
Zone replacer share			-0.037***	
			(0.001)	
Annual zone releaser share				0.103***
				(0.003)
Annual zone releaser share				-0.040***
				(0.001)
Time FE	Yes	Yes	Yes	Yes
Pipeline FE	Yes	Yes	Yes	Yes
Delivery and Receipt point FE	Yes	Yes	Yes	Yes
N	86,112	86,082	78,778	78,657
adj. R^2	0.777	0.773	0.726	0.721

Standard errors in parentheses
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The results in Table 4.13 suggest a strong relationship between market share and tariffs charged, both for replacers and releasers, suggestive of price-setting behavior. Moreover, this relationship is twice as large for releasers. Remembering that releasers benefit from higher tariffs (supply) and replacers benefit from lower tariffs (demand), a 1% increase in releaser market share is associated with a 0.09% increase in tariffs, whereas a 1% increase in replacer market share is

[9] We also tried including shipper fixed effects, but there is not enough remaining variation (e.g., only over time within shippers) because market shares tend to be fairly consistent over time.

associated with a 0.04% decrease in tariffs.

These results suggest two-sided market power. In such a context, it is unclear what the welfare implications of partially deregulating the natural gas transportation industry through its secondary market are. On the one hand, the secondary market allows tariffs to vary over time and respond to fluctuations in supply and demand. These tariffs are otherwise constant and set by long-term contracts in the primary market, which may prevent the efficient allocation of capacity, particularly during shocks. On the other hand, Deregulation of pipeline capacity implies that market participants influence tariffs based on their market share, as seen in Table 4.13, which may restrict the allocation of capacity. Thus, the aggregate implications are unclear, and there may be significant variation across pipelines and regions. We plan to investigate these welfare implications in future research.

4.5 Conclusion

This paper investigates a major deregulation effort in a highly regulated industry characterized by natural monopolies: the secondary market for natural gas capacity release. It is one of the first empirical investigations of a real Coasian market in which market participants can trade legal entitlements, allowing us to shed light on the resilience of supply chains, particularly on the supply of natural gas during periods of high demand. Our results highlight that the secondary market is an important market constituent during shocks, as evidenced by the large and persistent increase in capacity released in that market during the February 2021 cold wave in Texas. We argue that the secondary market improves the efficiency of capacity allocation by deregulating transportation tariffs and providing information to market participants, which fosters the formation of long-term relationships between releasers and replacers. However, deregulation comes at a cost, and we find evidence of market power that reflects price-setting behavior from both releasers and replacers.

Bibliography

Avalos, Roger, Timothy Fitzgerald, and Randal R. Rucker, "Measuring the effects of natural gas pipeline constraints on regional pricing and market integration," *Energy Economics*, 2016, *60.*

Bathiany, Sebastian et al., "Climate Models Predict Increasing Temperature Variability in Poor Countries," *Sciences Advances*, 2018, *4.*

Berger, David W., Kyle F. Herkenhoff, Andreas R. Kostøl, and Simon Mongey, "An Anatomy of Monopsony: Search Frictions, Amenities and Bargaining in Concentrated Markets," *NBER Working Paper No. 31149*, 2023.

Brown, Stephen P.A. and Mine K. Yücel, "Deliverability and regional pricing in U.S. natural gas markets," *Energy Economics*, 2008, *30.*

Chen, Yanyou, "Network Structure and Efficiency Gains from Mergers: Evidence from U.S. Freight Railroads," 2023.

Freyaldenhoven, Simon, Christian Hansen, and Jesse M. Shapiro, "Pre-Event Trends in the Panel Event-Study Design," *American Economics Review*, 2019, *109* (9).

_ , _ , Jorge Pérez Pérez, and Jesse M. Shapiro, "Visualization, identification, and estimation in the linear panel event-study design," *Advances in Economics and Econometrics: Twelfth World Congress*, 2024.

Gehrig, Thomas, "Intermediation in Search Markets," *Journal of Economics & Management Strategy*, 1993, *2.*

Gonçalves, Ana C. R., Xurxo Costoya, Raquel Nieto, and Margarida L. R. Liberato, "Extreme weather events on energy systems: A comprehensive review on impacts, mitigation, and adaptation measures," *Sustainable Energy Research*, 2024, *11.*

Harris, Adam and Thi Mai Anh Nguyen, "Long-Term Relationships and the Spot Market: Evidence from US Trucking," 2023.

Jasiūnas, Justinas, Peter D. Lund, and Jani Mikkola, "Energy system resilience – A review," *Renewable and Sustainable Energy Reviews*, 2021, *150.*

Marmer, Vadim, Dmitry Shapiro, and Paul MacAvoy, "Bottlenecks in regional markets for natural gas transmission services," *Energy Economics*, 2007, *29.*

Miller, Douglas L., "An Introductory Guide to Event Study Models," *Journal of Economic Perspectives*, 2023, *37* (2).

Oliver, Matthew E., Charles F. Mason, and David Finnoff, "Pipeline congestion and basis differentials," *Journal of Regulatory Economics*, 2014, *46.*

Salz, Tobias, "Intermediation and Competition in Search Markets: An Empirical Case Study," *Journal of Political Economy*, 2022, *130* (2).

Spulber, Daniel F., "Market Microstructure and Intermediation," *Journal of Economic Perspectives*, 1996, *10* (3).

Wang, Bin, Michela Biasutti, Michael P. Byrne, Christopher Castro, Chih-Pei Chang, Kerry Cook, Rong Fu, Alice MGrimm, Kyung Jaha, Harry Hendon et al., "Monsoons climate change assessment," *Bulletin of the American Meteorological Society*, 2021, *102.*

Waseem, Muhammad and Saeed D. Manshadi, "Electricity grid resilience amid various natural disasters: challenges and solutions," *Electricity Journal*, 2020, *33.*

Appendices

Appendix to Chapter 2

A Data

A.1 Details on Sampling Rules

In the ASI, Manufacturing plants are surveyed either as part of a census or as part of a sample. All plants who qualify for the census are required to fill out the survey by the Government of India's Central Statistics Office. The remaining plants are surveyed based on stratified sampling rules. The definition of census vs. sample and the sampling rules went through some changes over the years. In 2008, all plants with more than 100 workers and multi-plant firms, as well as plants in the lesser industrialized states (Manipur, Meghalaya, Nagaland, Tripura, Sikkim, and Andaman Nicobar Islands), were part of the census. Strata were constructed by state/industry pairs for the remaining plants, and 20% of plants were sampled within each stratum.

By 2016, the rules for a plant to be considered in the census expanded. Plants in the following states with more than 75 workers were part of the census: Jammu Kashmir, Himachal Pradesh, Rajasthan, Bihar, Chhattisgarh, and Kerala. Plants in the following states with more than 50 workers were part of the census: Chandigarh, Delhi, and Puducherry. Plants in the seven less industrialized states were part of the census: Arunachal Pradesh, Manipur, Meghalaya, Nagaland, Sikkim, Tripura, and Andaman Nicobar Islands. Lastly, the census included plants with more than 100 workers in all other states.

A.2 Calculating Emissions

To get establishment-level measures of greenhouse gas emissions, I convert units of potential energy (mmBtu) of each fuel into metric tons of carbon dioxide equivalent (CO_{2e}) as a result of combustion. Each mmBtu of fuel releases some quantity of carbon dioxide CO_2, methane CH_4, and nitrous oxide N_2O in the air, which may vary by industry based on standard practices and technology. Emissions of chemical k for a plant in industry j can be calculated as follows:

$$emissions_{jk} = \sum_f \sum_k \zeta_{fkj} * e_f$$

$$\forall\ k = \{CO_2, CH_4, N_2O\} \quad \forall\ f = \{\text{Natural Gas, Coal, Oil, Electricity}\}$$

The fuel-by-industry emission factors of each chemical ζ_{fkj} are found in the database provided by GHG Platform India and come from two main sources: India's Second Biennial Update Report (BUR) to United Nations Framework Convention on Climate Change (UNFCCC) and IPCC Guidelines. Quantities in mmBtu of each fuel e_f are observed for each establishment in each year. Then, quantities of each chemical are converted into carbon dioxide equivalent CO_2e using the Global Warming Potential (GWP) method as follows:

$$CO_{2e} = \underbrace{GWP_{CO_2}}_{=1} * CO_2 + GWP_{ch4} * CH_4 + GWP_{n2o} * N_2O$$

From the calculations above, I can define fuel-specific emission factors that will be used to directly convert fuels to CO_{2e} (or GHG). For fuel f in industry j (excluding electricity),

$$\gamma_{fj} = GWP_{co2} * \zeta_{f,co2,j} + GWP_{ch4} * \zeta_{f,ch4,j} + GWP_{n2o} * \zeta_{f,n2o,j}$$

Calculations of emissions from electricity are done slightly differently than from fossil fuels because emissions come from production rather than end usage of electricity. Figure A.1 shows that coal is used to consistently generate above 60% of total electricity in India, which increased in 2010 and started to decrease after 2012.

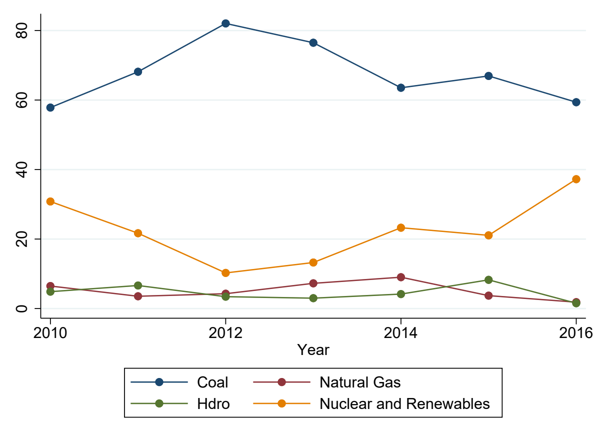

Figure A.1: Annual Indian Electricity Generation by Source (% of Total)
Source: International Energy Agency (IEA)

To construct measures of emissions from electricity, I take the distribution of emissions from different fuels used to produce electricity, averaged across years for the entire grid. Let $\omega_{ef} \in$

$[0, 1]$ $\forall f \in \{Coal, Gas\}$ be the share of fuel f used to generate electricity across the country, then

$$\gamma_{ej} = \sum_{f \in \{coal,gas\}} \omega_{ef} * \gamma_{fj}$$

Where γ_{fj} is the emission intensity of fuel f and was defined above.Total GHG emissions for plant i in industry j and year t is then defined as:

$$GHG_{ijt} = \gamma_e * e_{eijt} + \sum_{f \in \{natgas,coal,oil\}} \gamma_{fj} * e_{fijt}$$

Below are the tables detailing emissions factors. Note that for oil, I take the average over all petroleum fuels. The dispersion between oil types is much lower than the dispersion between the average of oil and coal/gas.

		Emission factors (kg CO_2e/mmBtu)			
Fuel	*Industry*	CO_2	CH_4	N_2O	Total (γ_{fj})
	Cement	100.90	0.03	0.42	101.34
	Non-ferrous metals	101.67	0.03	0.42	102.11
Coal	Pulp and paper	101.59	0.03	0.42	102.04
	Electricity generation	102.09	0.03	0.42	102.54
	Other	98.84	0.03	0.42	99.29
Oil	All	77.34	0.09	0.17	77.59
Natural Gas	All	50.64	0.03	0.03	50.70

Table A.1: Emission Factors From Fuels to Carbon Dioxide Equivalent

Notes: Emission factors are defined as $\zeta_{fkj} * GWP_k$ (kg CO_2e/mmBtu). Source: (**?**, Annexure 3)

Share of Electricity Generated by Source				
Natural Gas	Coal	Hydro	Other	Emission factor (kg CO_2e/mmBtu)
0.052	0.68	0.046	0.23	72.05

Table A.2: Emission Factors From Electricity

Steel ProductFuel Set	oil, elec	oil, elec, coal	oil, elec, gas	oil, elec, gas, coal	other	*Total*
Pig iron	50.46	23.39	8.29	5.71	12.15	100
Direct reduced iron	62.42	22.53	4.89	2.31	7.85	100
Ingots	54.12	17.51	9.23	6.76	12.38	100
Ferro-alloy	51.42	24.05	4.91	4.11	15.51	100
Hot and cold-rolled steel	38.32	21.14	15.00	14.80	10.75	100
Tubes	67.97	4.74	15.32	4.46	7.52	100
Wires	70.04	4.61	9.93	1.42	14.01	100
Other	47.59	19.79	10.35	9.96	12.30	100

Table A.3: Distribution of Fuels Sets by Steel Variety

A.3 Additional Evidence

B Model

B.1 Closing the Model: Aggregation Details

Given a mass of N_t operating plants, income I_t and aggregate demand shock e^{Γ_t}, the representative consumer solves:

$$\begin{aligned} \max_{\{Y_{it}\}_{i=1}^{N_t}, Y_{0t}} \mathbb{U} &= Y_{0t} + \frac{e^{\Gamma_t}}{\theta}\left(\frac{1}{N_t}\int_{\Omega_i}(N_t Y_{it})^{\frac{\rho-1}{\rho}}di\right)^{\frac{\theta\rho}{\rho-1}} \\ s.t. \quad Y_{0t} &+ \int_{\Omega_i} P_{it}Y_{it}di \leq I_t \end{aligned} \tag{4}$$

Following **?**, this can be separated into two problems. First, the consumers choose consumption of the aggregate final good Y_t, given some aggregate price index P_t and aggregate demand shock e^{Γ_t}:

$$\begin{aligned} &\max_{Y_{0t}, Y_t} Y_{0t} + \frac{e^{\Gamma_t}}{\theta}Y_t^{\theta} \\ &s.t. \quad Y_{0t} + P_t Y_t \leq I_t \end{aligned}$$

The optimal consumption of the aggregate final good is given by $Y_t(P_t) = \left(\frac{P_t}{e^{\Gamma_t}}\right)^{\frac{-1}{1-\theta}}$, and consumption of the outside good is given by $Y_{0t}(P_t) = I_t = P_tY_t(P_t) = I_t - e^{\Gamma_t\frac{1}{1-\theta}}P_t^{\frac{-\theta}{1-\theta}}$. Putting the two together yields the indirect utility $\mathbb{V}$, which corresponds to the consumer surplus due to quasi-linear preferences:

$$\mathbb{V} = I_t + \Big(\frac{1}{1-\theta}\Big)\Gamma_t^{\frac{1}{1-\theta}} P_t^{\frac{-\theta}{1-\theta}}$$

This is the same indirect utility function as in **?**, augmented with an aggregate demand shock. Consumer surplus is decreasing in the aggregate price index, keeping income constant. Then, the representative consumer chooses which varieties to allocate for a given quantity of good Y_t by minimizing the cost of different varieties:

$$\min_{\{Y_{it}\}_{i=1}^{N_t}} \int_{\Omega_i} P_{it} Y_{it} \quad s.t. \qquad Y_t = \left(\frac{1}{N_t}\int_{\Omega_i} (N_t Y_{it})^{\frac{\rho-1}{\rho}} di\right)^{\frac{\rho}{\rho-1}}$$

Solving this cost-minimization problem yields the following conditional demand for each variety:

$$Y_{it}(Y_t) = \frac{Y_t}{N_t}\Big(\frac{P_{it}}{P_t}\Big)^{-\rho} \tag{5}$$

Combining both steps together yields the demand for each variety, corresponding to equation 2.4 in the main text:

$$Y_{it} = \frac{e^{\Gamma_t \frac{1}{1-\theta}}}{N_t} P_t^{\frac{\rho(1-\theta)-1}{1-\theta}} P_{it}^{-\rho}$$

Where the aggregate price index is such that $\int_{\Omega_t} P_{it} Y_{it} di = P_t Y_t$ and is given by $P_t = \left(\frac{1}{N_t}\int_{\Omega_i} P_{it}^{1-\rho}\right)^{\frac{1}{1-\rho}}$.

C Identification

C.1 Derivation of Estimating Equation for Outer Production Function

Production function:

$$\frac{Y_{it}}{\overline{Y}} = e^{\omega_{it}}\left(\alpha_k\Big(\frac{K_{it}}{\overline{K}}\Big)^{\frac{\sigma-1}{\sigma}} + \alpha_l\Big(\frac{L_{it}}{\overline{L}}\Big)^{\frac{\sigma-1}{\sigma}} + \alpha_m\Big(\frac{M_{it}}{\overline{M}}\Big)^{\frac{\sigma-1}{\sigma}} + \alpha_e\Big(\frac{E_{it}}{\overline{E}}\Big)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\eta\sigma}{\sigma-1}} \tag{6}$$

$$= e^{\omega_{it}}\Big(\alpha_k\tilde{K}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_l\tilde{L}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_m\tilde{M}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_e\tilde{E}_{it}^{\frac{\sigma-1}{\sigma}}\Big)^{\frac{\eta\sigma}{\sigma-1}} \tag{7}$$

Where I define $\frac{X_{it}}{\overline{X}} = \tilde{X}_{it}$

L_{it}, M_{it}, E_{it} are flexible inputs

I observe the quantity for L_{it} and K_{it} but only spending for materials and energy: $S_{M_{it}}, S_{E_{it}}$

Profit-maximization subject to technology and demand constraint:

$$\begin{aligned}
&\max_{L_{it},M_{it},E_{it}} \Big\{P_{it}(Y_{it})Y_{it} - p_{Mit}M_{it} - p_{Eit}E_{it} - w_tL_{it}\Big\}\\
s.t.\quad & Y_{it} = \overline{Y}e^{\omega_{it}}\Big(\alpha_K\tilde{K}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_L\tilde{L}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_M\tilde{M}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_E\tilde{E}_{it}^{\frac{\sigma-1}{\sigma}}\Big)^{\frac{\eta\sigma}{\sigma-1}}\\
& P_{it}(Y_{it}) = \left(\frac{e^{\Gamma_t}}{N_tY_{it}}\right)^{\frac{1}{\rho}} P_t^{\frac{1+\rho(\theta-1)}{(\theta-1)\rho}}
\end{aligned}$$

First-order conditions:

M_{it}/L_{it}:

$$\frac{M_{it}}{\overline{M}} = \Big(\frac{\alpha_L}{\alpha_M}\frac{S_{Mit}}{S_{Lit}}\Big)^{\frac{\sigma}{\sigma-1}}\frac{L_{it}}{\overline{L}} \tag{8}$$

E_{it}/L_{it}:

$$\frac{E_{it}}{\overline{E}} = \Big(\frac{\alpha_L}{\alpha_E}\frac{S_{Eit}}{S_{Lit}}\Big)^{\frac{\sigma}{\sigma-1}}\frac{L_{it}}{\overline{L}} \tag{9}$$

L_{it}:

$$\Big(\frac{e^{\Gamma_t}}{N_t}\Big)^{\frac{1}{\rho}} P_t^{\frac{\rho(1-\theta)-1}{(1-\theta)\rho}} \frac{\rho-1}{\rho}\eta(e^{\omega_{it}}\overline{Y})^{\frac{\rho-1}{\rho}}\alpha_L L_{it}^{\frac{\sigma-1}{\sigma}} ces_{it}^{\frac{\rho[\sigma(\eta-1)+1]-\eta\sigma}{(\sigma-1)\rho}} = S_{L_{it}}$$

Where $ces_{it} = \Big(\alpha_k\tilde{K}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_l\tilde{L}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_m\tilde{M}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_e\tilde{E}_{it}^{\frac{\sigma-1}{\sigma}}\Big)$

Using the FOC for labor, I can solve for total factor productivity $e^{\omega_{it}}$:

$$e^{\omega_{it}\frac{\rho-1}{\rho}} = \overline{Y}^{\frac{\rho-1}{\rho}} \frac{\rho}{\rho-1}\frac{1}{\eta}\Big(\frac{N_t}{e^{\Gamma_t}}\Big)^{\frac{1}{\rho}} P_t^{\frac{1-\rho(1-\theta)}{(1-\theta)\rho}} \frac{S_{L_{it}}}{\alpha_L L_{it}^{\frac{\sigma-1}{\sigma}}} ces_{it}^{\frac{\eta\sigma-\rho[\sigma(\eta-1)+1]}{(\sigma-1)\rho}} \tag{10}$$

Plug (10) into revenue equation:

$$\begin{aligned}
R_{it} &= P_{it}(Y_{it})Y_{it}e^{u_{it}} \\
&= \Big(\frac{e^{\Gamma_t}}{N_t}\Big)^{\frac{1}{\rho}} P_t^{\frac{1+\rho(\theta-1)}{(\theta-1)\rho}} Y_{it}^{\frac{\rho-1}{\rho}} e^{u_{it}} \\
&= \Big(\frac{e^{\Gamma_t}}{N_t}\Big)^{\frac{1}{\rho}} P_t^{\frac{1+\rho(\theta-1)}{(\theta-1)\rho}} \Bigg(e^{\omega_{it}}\Big(\alpha_K\tilde{K}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_L\tilde{L}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_M\tilde{M}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_E\tilde{E}_{it}^{\frac{\sigma-1}{\sigma}}\Big)^{\frac{\eta\sigma}{\sigma-1}}\Bigg)^{\frac{\rho-1}{\rho}} e^{u_{it}} \\
&= \frac{\rho}{\rho-1}\frac{1}{\eta}\Big(\alpha_K\tilde{K}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_L\tilde{L}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_M\tilde{M}_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_E\tilde{E}_{it}^{\frac{\sigma-1}{\sigma}}\Big)e^{u_{it}}
\end{aligned}$$

Plug ratio of FOCs (8) and (9) into the previous equation:

$$\begin{aligned}
R_{it} &= \frac{\rho}{\rho-1}\frac{1}{\eta}S_{Lit}\Bigg(\frac{\alpha_k}{\alpha_L}\Big(\frac{\tilde{K}_{it}}{L_{it}}\Big)^{\frac{\sigma-1}{\sigma}} + 1 + \frac{S_{Mit}}{S_{Lit}} + \frac{S_{Eit}}{S_{Lit}}\Bigg)e^{u_{it}} \\
&= \frac{\rho}{\rho-1}\frac{1}{\eta}\Bigg(S_{Lit}\Big(1 + \frac{\alpha_k}{\alpha_L}\Big(\frac{\tilde{K}_{it}}{\tilde{L}_{it}}\Big)^{\frac{\sigma-1}{\sigma}}\Big) + S_{Mit} + S_{Eit}\Bigg)e^{u_{it}}
\end{aligned}$$

Estimating Equation:

$$\ln R_{it} = \ln\frac{\rho}{\rho-1} + \ln\frac{1}{\eta} + \ln\Bigg(S_{Lit}\Big(1 + \frac{\alpha_k}{\alpha_L}\Big(\frac{\tilde{K}_{it}}{\tilde{L}_{it}}\Big)^{\frac{\sigma-1}{\sigma}}\Big) + S_{Mit} + S_{Eit}\Bigg) + u_{it} \tag{11}$$

C.2 Computational Details on Solving the Dynamic Discrete Choice Model

I show how to iterate over the expected value function $\vec{W}$ until $\| \vec{W}^{n+1} - \vec{W}^{n} \|$ is small enough with a very large state space, where for any set of states today $s, \mathcal{F}$.

$$W^n(s,\mathcal{F}) = \gamma + \log\left(\sum_{\mathcal{F}'\in\mathbb{F}} \exp\left(\pi(s,\mathcal{F}) + \Phi(\mathcal{F}' \mid \mathcal{F}) + \beta \int W^n(s',\mathcal{F}')dF(s' \mid s)\right)\right)$$

To evaluate the expected value function, note that there are originally 12 state variables: prices and productivity of all four fuels, hicks neutral productivity, the price of material inputs, year of observation, and whether a plan is located near a pipeline. I can reduce the dimension of the state space to 8 state variables, 2 of which are deterministic and 6 of which follow a Markov process. The 6 Markovian state variables are hicks-neutral productivity z, price of materials p_m, price/productivity of electricity p_e/ψ_e, price/productivity of oil p_o/ψ_o, price/productivity of gas p_g/ψ_g, and price/productivity of coal p_c/ψ_c, which are allowed to be correlated. Then,

$$\begin{aligned}\int W^{n+1}(s',\mathcal{F}')dF(s' \mid s) = \int_z \int_{p_m} \int_{\frac{p_e}{\psi_e}} \int_{\frac{p_o}{\psi_o}} \int_{\frac{p_g}{\psi_g}} \int_{\frac{p_c}{\psi_c}} W^n\left(z', p'_m \frac{p'_e}{\psi'_e}, \frac{p'_o}{\psi'_o}, \frac{p'_g}{\psi'_g}, \frac{p'_c}{\psi'_c}, \mathcal{F}', t, d\right) \times \\ f_{z',p'_m,\frac{p'_e}{\psi'_e},\frac{p'_o}{\psi'_o},\frac{p'_g}{\psi'_g},\frac{p'_c}{\psi'_c}}\left(z', p'_m, \frac{p'_e}{\psi'_e}, \frac{p'_o}{\psi'_o}, \frac{p'_g}{\psi'_g}, \frac{p'_c}{\psi'_c} \middle| z, p_m, \frac{p_e}{\psi_e}, \frac{p_o}{\psi_o}, \frac{p_g}{\psi_g}, \frac{p_c}{\psi_c}\right) d_z d_{p_m} d_{\frac{p_e}{\psi_e}} d_{\frac{p_o}{\psi_o}} d_{\frac{p_g}{\psi_g}} d_{\frac{p_c}{\psi_c}}\end{aligned}$$

Where t corresponds to the year of observation and d is an indicator of access to a natural gas pipeline. I approximate this high dimensional expected value function by discretizing the state space and the underlying Markov process. Since most state variables are highly persistent AR(1) processes, I use Rouwenhorst (1995) to discretize the process. Let M be the number of points on each grid. I am currently using $M = 4$, which gives me $4^6 = 4,096$ grid points for the Markovian state variables. When adding the six years of observations between 2010 and 2015 as well as the access to the pipeline indicator, I get $(4^6)*6*2 = 49,152$. However, the curse of dimensionality starts to kick in when I add the distribution of comparative advantages for gas and coal (see later sections). With three possible values for gas and coal, this gives me nine possible combinations of comparative advantages. Ultimately, I am left with $(4^6)*6*2*9 = 442,368$ grid points. Using this discretization process, I can then represent the value function as a block matrix $\vec{W}^n$ containing all state combinations. Let S be the set of all state variable combinations, $\Gamma(s' \mid s)$ be the vector of all state transition probabilities when starting at state s (in vectorized form), $\mathbf{\Pi}$ be the vector of all possible profit combinations, $\vec{\mathcal{K}}$ be the vector of all possible fuel set switching costs. Then

$$\vec{W} \approx \gamma + \log\left(\sum_{\mathcal{F}'\in\mathbb{F}} \exp\left(\mathbf{\Pi} + \vec{\mathcal{K}}(\mathcal{F}') + \beta\Big[\bigotimes_{s\in S} \Gamma(s' \mid s)\Big]^T \vec{W}\right)\right) \tag{12}$$

Lastly, to reduce the computational burden, I iterate over equation (12) by parallelizing across all possible combinations of starting states using graphics processing units (GPU) Arrays with CUDA. Computational gains using GPU Arrays are significant over standard CPU parallelization. Detailed Julia code is available on my GitHub[10].

C.3 Details of Em Algorithm to Recover Distribution of Fixed Costs and Comparative Advantages

The procedure to estimate the fixed costs parameters θ_1 and the unselected, unconditional distribution of fuel-specific random effects is explained below. I experimented with both the **?** version that relies on a nested fixed point algorithm to update the value function and the **?** that uses the conditional choice probabilities (CCP) and forward simulations to update the value function. In the main version of the paper, I am using the nested fixed point version with a large grid for the state space as discussed in Appendix C.2.

$$\begin{aligned}
\ln \mathcal{L}(\mathcal{F}, s \mid \theta_1, \theta_2) &= \sum_{i=1}^{n} \ln\left[\sum_k \pi_k \left[\prod_{t=1}^{T} Pr(\mathcal{F}_{it+1} \mid \mathcal{F}_{it}, s_{it}, \mu_{fi} = \mu_k; \theta_1, \theta_2)\right]\right] + \sum_{i=1}^{n}\sum_{t=1}^{T} \ln f(s_{it} \mid s_{it-1}; \theta_2) \\
&= \sum_{i=1}^{n} \ln\left[\sum_k \pi_k \left(\prod_{t=1}^{T} \frac{e^{\upsilon_{\mathcal{F}_{it+1}}(\mathcal{F}_{it}, s_{it}, \mu_i = \mu_k; \theta_1, \theta_2)}}{\sum_{\mathcal{F}\subseteq\mathbb{F}} e^{\upsilon_{\mathcal{F}}(\mathcal{F}_{it}, s_{it}, \mu_i = \mu_k; \theta_1, \theta_2)}}\right)\right] + \sum_{i=1}^{n}\sum_{t=1}^{T} \ln f(s_{it} \mid s_{it-1}; \theta_2)
\end{aligned}$$

In principle, one can directly estimate both the fixed costs θ_1 and the distribution of comparative advantages from the full information likelihood above. However, this is computationally very expensive and rarely used in practice. For this reason, **?** use Baye's law to show that the first-order conditions of the full information likelihood with respect to all parameters are the same as the first-order conditions of the posterior likelihood with respect to fixed costs θ_1 given some prior guess of the distribution of unobserved heterogeneity. This is the key result that allows me

[10] `https://github.com/emmanuelmurrayleclair/JMP`

to use the EM algorithm.

$$\begin{aligned}\hat{\theta}_1 &= \underset{\theta_1,\theta_2,\pi}{\arg\max} \sum_{i=1}^{n} \ln\Bigg[\sum_k \pi_k \Bigg[\prod_{t=1}^{T} Pr(\mathcal{F}_{it+1} \mid \mathcal{F}_{it} s_{it}, \mu_i = \mu_k; \theta_1, \theta_2)\Bigg]\Bigg] \\ &\equiv \underset{\theta_1}{\arg\max} \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_k \rho(\mu_k \mid \mathcal{F}_i, s_i; \hat{\theta}_1, \hat{\theta}_2, \hat{\pi}) \ln Pr(\mathcal{F}_{it+1} \mid \mathcal{F}_{it}, s_{it}, \mu_i = \mu_k; \theta_1, \hat{\theta}_2)\end{aligned}$$

Estimation is iterative and proceeds as follows:

1. Estimate the distribution of state variables externally $\hat{\theta}_2$. These stay fixed throughout the procedure.
2. Initialize fixed cost parameters θ_1^0 and guess some initial probabilities $\{\pi_{f1}^0, \pi_2^0, ..., \pi_K^0\}$. I use the distribution of selected random effects to initialize this distribution.
3. Do value function iteration (VFI) to update the expected value function W for all combinations of state variables conditional on these guesses, where different realizations of the random effects μ_k are just another state variable that is fixed over time.

$$\begin{aligned}W(s, \mathcal{F}, \mu_k; \theta_1^0, \hat{\theta}_2) &= \gamma \\ &+ \log\Bigg(\sum_{\mathcal{F}' \in \mathbb{F}} \exp\Big(\pi(s, \mathcal{F}) + \mathcal{K}(\mathcal{F}' \mid \mathcal{F}, s; \theta_1^0) + \beta \int W(s', \mathcal{F}', \mu_k; \theta_1^0, \hat{\theta}_2) dF(s' \mid s; \hat{\theta}_2)\Big)\Bigg)\end{aligned}$$

4. Get posterior conditional probabilities that plant i is of type k, $\rho^1(\mu_k \mid \mathcal{F}_i, s_i; \theta_1^0, \hat{\theta}_2, \pi^0)$, according to Baye's law:

$$\rho^1(\mu_k \mid \mathcal{F}_i, s_i; \theta_1^0, \hat{\theta}_2, \pi^0) = \frac{\pi_{fk}^0 \Big[\prod_{t=1}^{T} \Big[\prod_{\mathcal{F} \subseteq \mathbb{F}} \Big[Pr(\mathcal{F}_{it} \mid s_{it}, \mu_i = \mu_k; \theta_1^0, \hat{\theta}_2)\Big]^{\mathbb{I}(\mathcal{F}_{it} = \mathcal{F})}\Big]\Big]}{\sum_k \pi_k^0 \Big[\prod_{t=1}^{T} \Big[\prod_{\mathcal{F} \subseteq \mathbb{F}} \Big[Pr(\mathcal{F}_{it} \mid s_{it}, \mu_i = \mu_k; \theta_1^0, \hat{\theta}_2)\Big]^{\mathbb{I}(\mathcal{F}_{it} = \mathcal{F})}\Big]\Big]}$$

5. **E-step:** Update the unconditional comparative advantage probabilities as follows:

$$\pi_k^1 = \frac{\sum_{i=1}^{n} \rho^1(\mu_k \mid \mathcal{F}_i, s_i; \theta_1^0, \hat{\theta}_2, \pi^0)}{n} \qquad \forall k$$

6. **M-step:** Find fixed cost parameters θ_1^1 that maximize the (log)-likelihood conditional on current guess of unconditional and conditional probabilities $\pi_k^1, \rho^1(\mu_k \mid .)$
7. Repeat 3-6 until the full information likelihood is minimized.

D Counterfactuals

D.1 Discretizing the Process for Fuel Prices and Productivity Separately

The problem to solve is that I need to separate fuel prices from fuel productivity when studying the impact of a per-unit carbon tax levied on fossil fuels $(p_{fit} + \tau_f)$ because $\frac{p_{fit}+\tau_f}{\psi_{fit}} = \frac{p_{fit}}{\psi_{fit}} + \frac{\tau_f}{\psi_{fit}}$

Initially, the model is estimated with a process for the log of fuel prices/productivity from Equation 2.18, which I discretize into a Markov Chain. The Markov chain is a sequence of fuel prices/productivity realizations $\ln \frac{p_{f1}}{\psi_{f1}}, \ln \frac{p_{f2}}{\psi_{f2}}, \ln \frac{p_{f3}}{\psi_{f3}}, ...$ such that,

$$Pr\Big(\ln \frac{p_{fk+1}}{\psi_{fk+1}}\Big| \ln \frac{p_{fk}}{\psi_{fk}}, \ln \frac{p_{fk-1}}{\psi_{fk-1}}, \ln \frac{p_{fk-2}}{\psi_{fk-2}}, ...\Big) = Pr\Big(\ln \frac{p_{fk+1}}{\psi_{fk+1}}\Big| \ln \frac{p_{fk}}{\psi_{fk}}\Big)$$
$$Pr\Big(\ln \frac{p_{fk+1}}{\psi_{fk+1}}\Big| \ln \frac{p_{fk}}{\psi_{fk}}\Big) > 0 \quad \forall \quad k$$

From this Markov chain for fuel price/productivity, I create two Markov Chains: one for prices $\ln p_{f1}, \ln p_{f2}, ...$ and one for productivity $\ln \psi_{f1}, \ln \psi_{f2}, ...$ such that,

$$\ln p_{fk} = \ln \frac{p_{fk}}{\psi_{fk}} + \ln \psi_{fk}$$
$$Pr\Big(\ln p_{fk+1} \mid \ln p_{fk}\Big) = Pr\Big(\ln \psi_{fk+1} \mid \ln \psi_{fk}\Big) = Pr\Big(\ln \frac{p_{fk+1}}{\psi_{fk+1}}\Big| \ln \frac{p_{fk}}{\psi_{fk}}\Big)$$

To find the grid points for fuel prices and fuel productivity, I match the moments of the newly constructed Markov chains with moments from the distribution of fuel prices, fuel productivity, and other state variables in the data. All the moments I use include the variance of fuel prices, the variance of fuel productivity, the covariance between fuel prices and fuel productivity, the covariance between fuel prices and all other states, and the covariance between fuel productivity and all other states.

D.2 Shapley-Owen-Shorrocks Decomposition of the Carbon Tax Revenues Along Subsidy Rate

Given a subsidy rate $s \in [0, 1]$ towards the fixed costs of natural gas, the function of interest is the change in tax revenues/externality damages between the economy with subsidy s and the economy with no subsidy:

$$\begin{aligned}\Delta\mathcal{T}(s) &= \mathcal{T}(s) - \mathcal{T}(0) \\ &= \mathbb{E}_0\Big(\sum_{t=0}\beta^t\sum_i \tau_f e_{fit}(s)\Big) - \mathbb{E}_0\Big(\sum_{t=0}\beta^t\sum_i \tau_f e_{fit}(0)\Big)\end{aligned}$$

Note that this function is implicitly a function of all state variables and the carbon tax rate. As such, it can be decomposed into multiple arguments. The two arguments of interest here are the total quantity of energy E, which corresponds to the scale effect (higher quantity of energy equals higher quantities of all fuels), and the distribution of fuel sets in the economy $\mathcal{F}$, which corresponds to the substitution effect. Then, expected tax revenues can be rewritten as:

$$\begin{aligned}\Delta\mathcal{T}(\underbrace{E(s)}_{\text{scale}}, \underbrace{\mathcal{F}(s)}_{\text{substitution}}) &= \mathcal{T}(E(s), \mathcal{F}(s)) - \mathcal{T}(E(0), \mathcal{F}(0)) \\ &= \mathbb{E}_0\Big(\sum_{t=0}\beta^t\sum_i \tau_f e_{fit}(E_{it}(s), \mathcal{F}_{it}(s))\Big) - \mathbb{E}_0\Big(\sum_{t=0}\beta^t\sum_i \tau_f e_{fit}(E_{it}(0), \mathcal{F}_{it}(0))\Big)\end{aligned}$$

Here, I define the null case for both scale $\emptyset_E$ and substitution $\emptyset_{\mathcal{F}}$ arguments as the case with no substitute, such that the function is well defined, satisfying the criteria laid out in **?**.

$$\Delta\mathcal{T}(\emptyset_E, \emptyset_{\mathcal{F}}) = \mathcal{T}(E(0), \mathcal{F}(0)) - \mathcal{T}(E(0), \mathcal{F}(0)) = 0$$

Since there are two arguments, there will always be only two submodels that can exclude each argument: when the other argument is present and when it is not, with an associated probability of $\frac{1}{2}$ for each submodel. It is then quite easy to show that the total partial effect of adding the scale and substitution effect, respectively, is as follows:

$$
\begin{aligned}
C_{\text{scale}} &= \frac{1}{2}\underbrace{\Big(\Delta\mathcal{T}(E(s),\mathcal{F}(s)) - \Delta\mathcal{T}(\emptyset_E,\mathcal{F}(s))\Big)}_{\text{adding scale with substitution}} + \frac{1}{2}\underbrace{\Big(\Delta\mathcal{T}(E(s),\emptyset_{\mathcal{F}}) - \Delta\mathcal{T}(\emptyset_E,\emptyset_{\mathcal{F}})\Big)}_{\text{adding scale without substitution}} \\
&= \frac{1}{2}\Big(\Delta\mathcal{T}(s) - \Delta\mathcal{T}(\emptyset_E,\mathcal{F}(s))\Big) + \frac{1}{2}\Big(\Delta\mathcal{T}(E(s),\emptyset_{\mathcal{F}})\Big)
\end{aligned}
$$

$$
\begin{aligned}
C_{\text{substitution}} &= \frac{1}{2}\underbrace{\Big(\Delta\mathcal{T}(E(s),\mathcal{F}(s)) - \Delta\mathcal{T}(E(s),\emptyset_{\mathcal{F}})\Big)}_{\text{adding substitution with scale}} + \frac{1}{2}\underbrace{\Big(\Delta\mathcal{T}(\emptyset_E,\mathcal{F}(s)) - \Delta\mathcal{T}(\emptyset_E,\emptyset_{\mathcal{F}})\Big)}_{\text{adding substitution without scale}} \\
&= \frac{1}{2}\Big(\Delta\mathcal{T}(s) - \Delta\mathcal{T}(E(s),\emptyset_{\mathcal{F}})\Big) + \frac{1}{2}\Big(\Delta\mathcal{T}(\emptyset_E,\mathcal{F}(s))\Big)
\end{aligned}
$$

Lastly, it can be seen that this decomposition satisfies the additive criteria laid out by **?**:

$$
\begin{aligned}
C_{\text{scale}} + C_{\text{substitution}} &= \Delta\mathcal{T}(s) - \frac{1}{2}\Delta\mathcal{T}(\emptyset_E,\mathcal{F}(s)) + \frac{1}{2}\Delta\mathcal{T}(\emptyset_E,\mathcal{F}(s)) + \frac{1}{2}\Delta\mathcal{T}(E(s),\emptyset_{\mathcal{F}}) - \frac{1}{2}\Delta\mathcal{T}(E(s),\emptyset_{\mathcal{F}}) \\
&= \Delta\mathcal{T}(s)
\end{aligned}
$$

D.3 Energy Production Function with Energy Productivity – Identification and Results

The energy production function is as follows:

$$
E_{it} = \psi_{Eit}\Big(\sum_{f\in\mathcal{F}_{it}} \beta_f e_{fit}^{\frac{\lambda-1}{\lambda}}\Big)^{\frac{\lambda}{\lambda-1}} \qquad \sum_{f\in\{o,g,c,e\}} \beta_f = 1
$$

Assuming that the log of energy productivity follows and AR(1) process with year dummies $\ln\psi_{Eit} = \mu_0^{\psi_E} + \mu_t^{\psi_E} + \rho_{\psi_E}\ln\psi_{Eit-1} + \epsilon_{it}^{\psi_E}$, the production function can be written in log as

$$
\ln E_{it} = \mu_0^{\psi_E} + \mu_t^{\psi_E} + \frac{\lambda}{\lambda-1}\Big(\sum_{f\in\mathcal{F}_{it}} \beta_f e_{fit}^{\frac{\lambda-1}{\lambda}}\Big) + \rho_{\psi_E}\ln E_{it-1} - \rho_{\psi_E}\frac{\lambda}{\lambda-1}\ln\Big(\sum_{f\in\mathcal{F}_{it}} \beta_f e_{fit-1}^{\frac{\lambda-1}{\lambda}}\Big) + \epsilon_{it}^{\psi_E}
$$

This is very similar to the estimating equation for the fully flexible energy production function in the main text, where $\epsilon_{it}^{\psi_E}$ is the innovation to energy productivity between $t-1$ and t. As such, it is independent of period $t-1$ decisions:

$$\mathbb{E}(\epsilon_{it}^{\psi_E} \mid \mathcal{I}_{it-1}) = 0$$

However, this innovation is correlated with fuel choices at t. I instrument fuel choices at t with aggregate variation in fuel prices due to exogenous reasons such as geopolitical events, which I interact with the share of each fuel to generate electricity by Indian States. These shift-share instruments are the same instruments proposed by **?**, which I also use to estimate demand in the main model. Together, these instruments and fuel choices at $t-1$ form a set of moment conditions that satisfy conditional independence of the error term and identify the relevant parameters of the production function: $\lambda, \beta_o, \beta_g, \beta_c, \beta_e$. Below are the estimates of the production function:

Table D.1: Estimates of Energy Production Function with Energy Productivity

	Steel	
Elasticity of substitution $\hat{\lambda}$	2.173***	(0.240)
Relative productivity of oil $\hat{\beta}_o$	0.099***	(0.011)
Relative productivity of gas $\hat{\beta}_g$	0.049***	(0.012)
Relative productivity of coal $\hat{\beta}_c$	0.426***	(0.033)
Observations	3459	

Standard errors in parentheses

$^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table D.2: Estimates of the Restricted Energy Production Function

Elasticity of the Price of Energy With Respect to Relative Fuel Prices

The price of energy in the fully flexible model is as follows:

$$p_{E_{it}} = \Big(\sum_{f \in \mathcal{F}_{it}} (p_{fit}/\psi_{fit})^{1-\lambda} \Big)^{\frac{1}{1-\lambda}}$$

and can be written in terms of price ratios for a given fuel (say gas), where $\tilde{p}_{fit} = p_{fit}/p_{git}$ and likewise for $\tilde{\psi}_{fit}$

$$p_{E_{it}} = p_{git}\Big(\sum_{f\in\mathcal{F}_{it}} (\tilde{p}_{fit}/\tilde{\psi}_{fit})^{1-\lambda}\Big)^{\frac{1}{1-\lambda}}$$

The the elasticity of this price of energy with respect to relative fuel prices (say coal relative to gas) is as follows:

$$\begin{aligned}
\frac{\partial \ln p_{E_{it}}}{\partial \ln(p_{cit}/p_{git})} &= \frac{1}{p_{E_{it}}}\Bigg[\frac{p_{git}}{1-\lambda}\Big(\sum_{f\in\mathcal{F}_{it}}(\tilde{p}_{fit}/\tilde{\psi}_{fit})\Big)^{\frac{\lambda}{\lambda-1}}\frac{\partial \exp((1-\lambda)(\ln\tilde{p}_{cit}-\ln\tilde{\psi}_{cit}))}{\partial \ln p_{fit}}\Bigg] \\
&= \frac{1}{p_{E_{it}}}\Bigg[p_{git}\Big(\sum_{f\in\mathcal{F}_{it}}(\tilde{p}_{fit}/\tilde{\psi}_{fit})\Big)^{\frac{\lambda}{\lambda-1}}(\tilde{p}_{cit}/\tilde{\psi}_{cit})^{1-\lambda}\Bigg] \\
&= \frac{(\tilde{p}_{cit}/\tilde{\psi}_{cit})^{1-\lambda}}{\sum_{f\in\mathcal{F}_{it}}(\tilde{p}_{fit}/\tilde{\psi}_{fit})^{1-\lambda}} = \frac{(p_{cit}/\psi_{cit})^{1-\lambda}}{\sum_{f\in\mathcal{F}_{it}}(p_{fit}/\psi_{fit})^{1-\lambda}}
\end{aligned}$$

Moreover, this elasticity is equal to the spending share of coal relative to all other fuels. To see this, relative first-order conditions of the cost-minimization problem in (2.5) for two fuels (c, g) are:

$$\begin{aligned}
\frac{p_{cit}}{p_{git}} &= \Big(\frac{\psi_{cit}e_{cit}}{\psi_{git}e_{git}}\Big)^{-\frac{1}{\lambda}}\frac{\psi_{cit}}{\psi_{git}} \\
\frac{e_{cit}}{e_{git}} &= \Big(\frac{p_{git}}{p_{cit}}\Big)^{\lambda}\Big(\frac{\psi_{cit}}{\psi_{git}}\Big)^{\lambda-1}
\end{aligned}$$

Multiplying both sides by relative prices, this yields:

$$\frac{p_{cit}e_{cit}}{p_{git}e_{git}} = \frac{(p_{cit}/\psi_{cit})^{1-\lambda}}{(p_{git}/\psi_{git})^{1-\lambda}}$$

Summing across all relative fuel spending shares yields the elasticity:

$$\frac{p_{cit}\psi_{cit}}{\sum_{f\in\mathcal{F}_{it}} p_{fit}e_{fit}} = \frac{(p_{cit}/\psi_{cit})^{1-\lambda}}{\sum_{f\in\mathcal{F}_{it}}(p_{fit}/\psi_{fit})^{1-\lambda}}$$

D.4 Correlation Between Fuel Productivity and Hicks-Neutral Productivity

	Total Factors	**Gas**	**Coal**	**Oil**	**Elec**
Total Factors	1				
Gas	-0.19	1			
Coal	-0.11	0.14	1		
Oil	-0.24	0.02	0.02	1	
elec	-0.56	0.13	0.07	0.06	1

Table D.3: Correlation Matrix of Fuel Productivity and Total Factor Productivity

D.5 Trade-off Across Values of Demand Elasticity

Below, I show the trade-off between aggregate output and emissions for different carbon tax levels, comparing how much better the economy fairs when allowing for heterogeneity in fuel productivity rather than just energy productivity. I do this exercise by varying the elasticity of demand, which affects the extent of output reallocation across plants. Consistent with **?**, I find that more elastic demand increases output reallocation across plants when allowing for fuel productivity, which increases the gap between the two production frontiers. This confirms the importance of the output reallocation channel in explaining the aggregate trade-off in the main text.

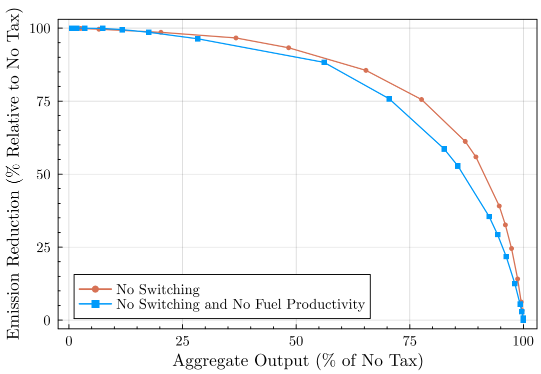

(a) Perfect Complement Output Varieties

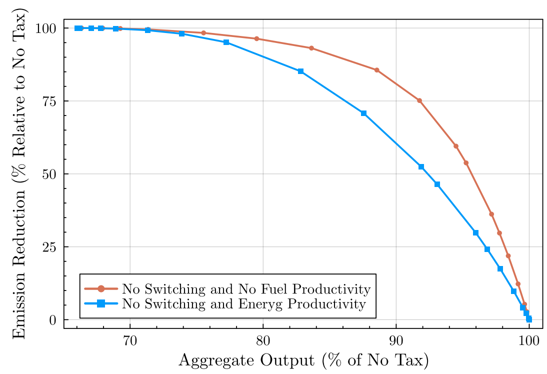

(b) Elasticity of Demand = 8

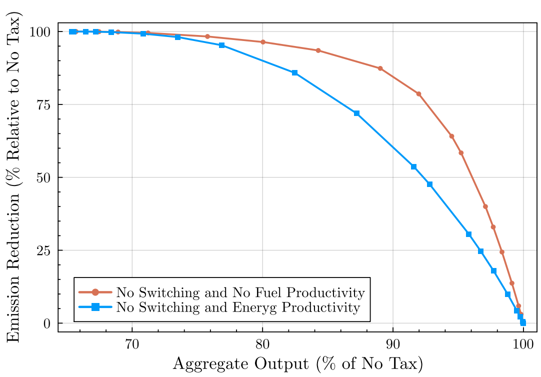

(c) Elasticity of Demand = 10

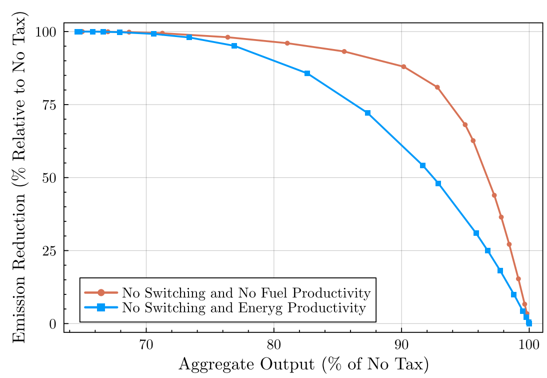

(d) Elasticity of Demand = 12

Figure D.1: Aggregate Trade-off Between Emission Reduction and Output, Across Different Elasticities of Demand

Notes: Figure (a) corresponds to perfect complement output varieties and refers to a representative consumer who has Leontief preferences across different varieties. While both production frontiers are closest to each other with such preferences, there is still a gap due to initial differences in fuel concentration. The economy with heterogeneity in fuel productivity allows for more concentration of fuels across plants, particularly coal. It thus yields more substitution away from coal in *level*, even though the elasticity is the same.

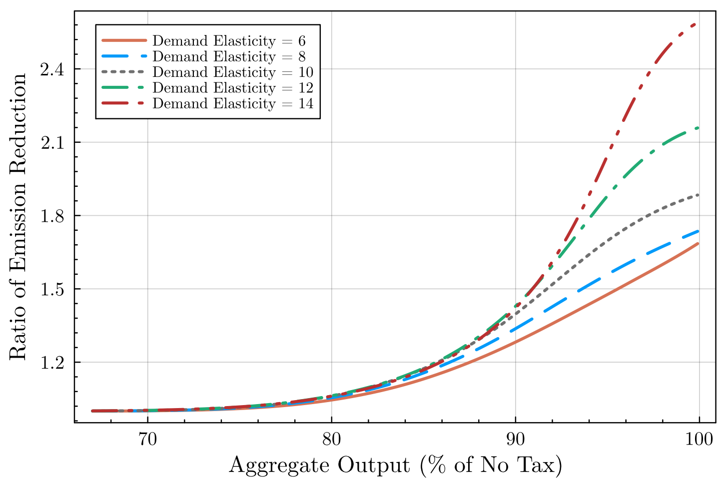

Figure D.2: Comparison of the Gap Between the Economy With Fuel Productivity and the Economy With Energy Productivity

Appendix to Chapter 3

E Preliminary Evidence for Entry/Exit

To investigate entry/exit as a result of the carbon tax, I run a simple difference-in-difference regression with the number of firms as the dependent variable. In both B.C. and Quebec, the control group comprises firms in all other provinces, and the treated period is from 2008 onwards for both regulated provinces. As the asymmetric carbon tax raises the marginal cost of regulated firms relative to the marginal cost of unregulated firms, standard Melitz theory suggests that the minimum productivity required to operate in regulated provinces would increase, decreasing the number of operating firms, and vice versa for unregulated provinces. Hence, the DiD coefficients should be negative. Here, it is positive and I cannot reconcile this finding with the theory, which is why I assume that the distribution of productivity remains the same after the tax.

	(1) British Columbia	(2) Quebec
time	-100.0***	-100.0***
	(2.57)	(2.43)
treated	-97.24***	-21.74***
	(4.19)	(3.38)
DiD	79.40***	71.66***
	(5.55)	(4.47)
constant	242.7***	242.7***
	(1.89)	(1.79)
N	7,888	8,813

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table E.1: Effect of Carbon Tax on Firm Entry and Exit

F Derivation Perceived Prices and Relative Fuel Quantities

Starting from the firm's cost minimization problem:

$$\min_{\{q_s^\ell\}_{\ell=1}^L} \Big\{ \sum_\ell p_{\ell s} q_s^\ell \Big\} \text{ s.t. } F = \Big(\sum_\ell \lambda_\ell (q_s^\ell)^{(\sigma-1)/\sigma} \Big)^{\sigma/(\sigma-1)}$$
$$\mathcal{L} = \sum_{\ell=1} p_{\ell s} q_s^\ell + \mu \Bigg(F - \Big(\sum_\ell \lambda_\ell (q_s^\ell)^{(\sigma-1)/\sigma} \Big)^{\sigma/(\sigma-1)} \Bigg)$$

FOC:

$$p_{\ell s} = \mu \Bigg(\frac{\lambda_\ell}{q_s^{\ell\,1/\sigma}} \Big(\sum_\ell \lambda_\ell q_s^{\ell(\sigma-1)/\sigma} \Big)^{1/(\sigma-1)} \Bigg) \,\forall\, \ell$$

I can divide fuel ℓ's FOC by fuel k's FOC:

$$\frac{p_{\ell s}}{p_{ks}} = \Big(\frac{q_s^k}{q_s^\ell} \Big)^{1/\sigma} \frac{\lambda_\ell}{\lambda_k}$$

Then,

$$q_s^\ell = \Big(\frac{p_{ks}}{p_{\ell s}} \frac{\lambda_\ell}{\lambda_k} \Big)^\sigma q_s^k$$

I can plug $q_s^\ell(q_s^k)$ into the technology:

$$F = \Bigg(\sum_\ell \Big[\Big(\frac{p_{ks}}{p_{\ell s}} \frac{\lambda_\ell}{\lambda_k} \Big)^\sigma q_s^k \Big]^{(\sigma-1)/\sigma} \Bigg)^{\sigma/(\sigma-1)}$$
$$= q_s^k \Big(\frac{p_{ks}}{\lambda_k} \Big)^\sigma \Big(\sum_\ell \lambda_\ell^\sigma p_{\ell s}^{1-\sigma} \Big)^{\sigma/(\sigma-1)}$$

Now I can define perceived prices:

$$\tilde{p}_{ks} = \frac{p_{ks}}{\lambda_k}$$
$$\tilde{p}_s = \Big(\sum_{\ell} \lambda_\ell^\sigma p_{\ell s}^{1-\sigma}\Big)^{1/(1-\sigma)}$$

Then, I get the optimal quantity of each fuel consumed, for a given output quantity F.

$$F = q_s^k \Big(\frac{\tilde{p}_{ks}}{\tilde{p}_s}\Big)^\sigma$$
$$q_s^k = \Big(\frac{\tilde{p}_{ks}}{\tilde{p}_s}\Big)^{-\sigma} F$$

G Derivation of the Log-Likelihood for Estimation of Macro Parameters

Recall the main estimating equation (eq. 10 in the paper):

$$\underbrace{\ln q_{ijst} + \hat{\sigma}(\ln \tilde{\mathbf{p}}_{\ell\mathbf{jst}} - \ln \tilde{\mathbf{p}}_{\mathbf{jst}})}_{\mathbf{y_{ijst}}} = \underbrace{\ln\Big(\frac{\rho - 1}{\rho}\Big) + \ln \beta_{jt} + \ln C_t - \rho \ln \tilde{p}_{jst} - \ln(\tilde{p}_{jrt}^{1-\rho} N_{jt}^r + \tilde{p}_{jurt}^{1-\rho} N_{jt}^{ur})}_{X_{ijst}(\rho)}$$
$$+ \underbrace{(\rho - 1)(\ln A_{ijst} - \ln \tilde{A}_{jt}) + \mathbf{u_{ijst}}}_{\epsilon_{ijst}}$$

Assuming that $g_{jt}(A) = g(A)$, then $\tilde{A}_{jt} = \tilde{A}$. I now assume a normal distribution for the distribution of log-productivity and both measurement errors:

$$A_{ijst} \sim LN(\mu, \sigma_A^2)$$
$$\iff \ln A_{ijst} \sim N(\mu, \sigma_A^2)$$
$$u_{ijst}^\ell \sim N(0, \sigma_\ell^2)$$

Recall the definition of aggregate productivity:

$$\tilde{A} = \Big(\int_0^\infty A^{\rho-1} g(A) dA \Big)^{1/(\rho-1)}$$

Using the moment-generating function for a log-normal distribution, I know that,

$$\mathrm{E}[A^{\rho-1}] = \exp\Big((\rho-1)\mu + (\rho-1)^2 \sigma_A^2/2 \Big)$$

I also know that $\tilde{A} = \mathrm{E}[A^{\rho-1}]^{1/(\rho-1)}$. Then,

$$\tilde{A} = \Big(\exp\Big(\mu(\rho-1) + (\rho-1)^2 \sigma_A^2/2 \Big) \Big)^{1/(\rho-1)}$$

I now have everything to derive the distribution of the relative log-productivity term where $\ln \tilde{A}$ is just a constant that gets added to the mean but does not change the variance:

$$(\rho-1)(\ln A_{ijst} - \ln \tilde{A}) \sim N\Big(-(\rho-1)^2 \sigma_A^2/2, (\rho-1)^2 \sigma_A^2 \Big)$$

I can now write the joint distribution for the composite error term ϵ_{ijst}, which will be used to form the log-likelihood:

$$\begin{pmatrix} \epsilon_{ijst}^{ng} \\ \epsilon_{ijst}^{o} \end{pmatrix} \sim N\left(\begin{pmatrix} -(\rho-1)^2 \sigma_A^2/2 \\ -(\rho-1)^2 \sigma_A^2/2 \end{pmatrix}, \begin{pmatrix} (\rho-1)^2 \sigma_A^2 + \sigma_{ng}^2 & (\rho-1)^2 \sigma_A^2 + \sigma_{ng,o} \\ (\rho-1)^2 \sigma_A^2 + \sigma_{ng,o} & (\rho-1)^2 \sigma_A^2 + \sigma_o^2 \end{pmatrix} \right)$$

H Derivation of Variance for Fuel Efficiency Parameters

Recalling the relative fuel equation to estimate technology parameters when I only have two fuels:

$$\ln q^{ng}_{ijst} - \ln q^{o}_{ijst} = \sigma(\ln p_{o,st} - \ln p_{ng,st}) + \Gamma_j + \epsilon_{ijst}$$

Where,

$$\Gamma_j = \sigma(\ln \lambda_{j,ng} - \ln \lambda_{j,o})$$

To estimate these reduced-form industry fixed effects, I redefine them as a constant plus an industry-specific term: $\Gamma_j = \gamma + \delta_j$. To get the variance of $\lambda_{j,o}, \lambda_{j,ng}$, I first need to derive the variance of $\Gamma_j = \gamma + \delta_j$ where γ and δ_j are the estimated parameters:

$$\sqrt{n}(\hat{\gamma} + \hat{\delta}_j - (\gamma + \delta_j)) \sim \mathcal{N}(0, \sigma^2_\gamma + \sigma^2_{\delta_j})\ \forall j$$
$$\hat{\sigma}^2_{\Gamma_j} = \hat{\sigma}^2_\gamma + \hat{\sigma}^2_{\delta_j}$$

Next, I can use the delta method to recover the variance of structural parameters, $\lambda_{ng,j}, \lambda_{oil,j}\ \forall j$, where I showed in the paper that those are functions of the interfuel elasticity of substitution, σ and reduced-form parameters Γ_j.

$$\hat{\lambda}_{ng,j} = \frac{\exp(\hat{\Gamma}_j/\hat{\sigma})}{\exp(\hat{\Gamma}_j/\hat{\sigma}) + 1}$$
$$\hat{\lambda}_{oil,j} = \frac{1}{\exp(\hat{\Gamma}_j/\hat{\sigma}) + 1}$$

By the delta method,

$$\sqrt{n}\begin{bmatrix} (\hat{\lambda}_{ng,j} - \lambda_{ng,j}) \\ (\hat{\lambda}_{oil,j} - \lambda_{oil,j}) \end{bmatrix} \sim \begin{bmatrix} \frac{\partial \lambda_{ng,j}}{\partial \Gamma} & \frac{\partial \lambda_{ng,j}}{\partial \sigma} \\ \frac{\partial \lambda_{oil,j}}{\partial \Gamma} & \frac{\partial \lambda_{oil,j}}{\partial \sigma} \end{bmatrix} \begin{bmatrix} \mathcal{N}(0, \sigma^2_{\Gamma_j}) \\ \mathcal{N}(0, \sigma^2_\sigma) \end{bmatrix}$$

Where the gradients are as follows:

$$\begin{bmatrix} \frac{\partial \lambda_{ng,j}}{\partial \Gamma} & \frac{\partial \lambda_{ng,j}}{\partial \sigma} \\ \frac{\partial \lambda_{oil,j}}{\partial \Gamma} & \frac{\partial \lambda_{oil,j}}{\partial \sigma} \end{bmatrix} = \begin{bmatrix} \frac{\exp(\Gamma_j/\sigma)}{\sigma(\exp(\Gamma_j/\sigma)+1)^2} & \frac{-\Gamma_j \exp(\Gamma_j/\sigma)}{\sigma^2(\exp(\Gamma_j/\sigma)+1)^2} \\ \frac{-\exp(\Gamma_j/\sigma)}{\sigma(\exp(\Gamma_j/\sigma)+1)^2} & \frac{\Gamma_j \exp(\Gamma_j/\sigma)}{\sigma^2(\exp(\Gamma_j/\sigma)+1)^2} \end{bmatrix}$$

Hence, the asymptotic variance of structural fuel shares by industry is as follows:

$$\begin{aligned} \sigma^2_{\lambda_{ng,j}} &= \Big(\frac{\exp(\Gamma_j/\sigma)}{\sigma(\exp(\Gamma_j/\sigma)+1)^2}\Big)^2 \sigma^2_{\Gamma_j} + \Big(\frac{\Gamma_j \exp(\Gamma_j/\sigma)}{\sigma^2(\exp(\Gamma_j/\sigma)+1)^2}\Big)^2 \sigma^2_{\sigma} \\ &= \sigma^2_{\lambda_{oil,j}} \end{aligned}$$

And I can use sample estimates of all those quantities to get the sample variance of fuel shares.

Appendix to Chapter 4

I Rate Regulation in the US Pipeline Industry

The United States pipeline regulatory authority, the Federal Energy Regulatory Commission (FERC), has an established framework and legislation allowing it to approve pipelines charging market-based rates. Much of this framework and legislation was developed in the late 80s and early 90s, and today, a handful of pipelines charge market-based rates, mostly oil pipelines. While the FERC grants oil and natural gas pipelines market-based rate authority under two separate pieces of legislation, the following must be included by a pipeline as evidence to move towards market-based rates:

- Describe the proposed service.
- Define relevant product and geographic markets.
- Provide the applicant's ownership and list affiliated energy companies, services provided, and their location. If an affiliate operates in the same geographic market, the market shares of the applicant and the affiliate should be combined.
- Identify good alternatives to the proposed service - which parties provide similar services within the same geographic market. List the applicant's competitors and location.
- Include market share and Herfindahl-Hirschman Index (HHI) calculations to measure market concentration. The Commission's traditional HHI threshold is 1,800.
- Discuss other relevant competitive factors, such as ease of entry and excess capacity held by competitors.
- Describe how the applicant's rates compare to the competitors.

In addition to the evidence above, some pipelines have also provided econometric analysis to show that their shippers have a relatively high elasticity of demand as a means to quantify the competitive position these pipelines face. It has also been required that pipelines provide cross-price elasticities for different products transported in their pipeline or by competitors. For example, a liquids pipeline may be constructed to ship various types of petroleum products (heavy crude, light crude, condensate, refined products, etc.) vs. another pipeline within the same area only capable of transporting a single type of petroleum product. In that case, the regulator does not see the product transported through the pipeline as homogenous and may review cross-price elasticities.

The FERC also has the authority to determine if only a certain portion of a system may be granted market-based rate authority. For example, the origin market for pipelines is often considered as the production areas, which often could be considered not sufficiently concentrated and may result in the pipeline having market power in the origin market. However, before reaching their ultimate destination, many pipelines pass through storage hubs. These storage hubs often have many pipelines entering them and are seen as a more competitive market. Therefore, a pipeline may not be granted market-based rates for its entire system but only a portion of its system, say, from these storage hubs to its destination market.

As is evident, the United States has a much more established framework for determining market-based rate authority. This is largely due to the more competitive nature of the pipeline industry in the United States compared to Canada. While this framework is established, it is also important to note that it represents a relatively steep bar to cross, and very few pipelines are granted this authority.

J Capacity Release Market

J.1 How Does Capacity Release Market Work?

We first explain in more detail the capacity release programs:

1. Capacity releases can be put forward by the pipeline or requested from shippers. Both get posted on the pipeline's website.
2. There are generally two types of capacity releases: biddable and non-biddable. Biddable capacity releases are subject to an open season (bidding process) where the highest bid gets the capacity. Non-biddable capacity releases are generally agreements that are entered into before posting on the regulatory website and are only eligible for releases with contract lengths of 31 days or less or more than one year. This is why we see most of the releases at 31 days in length or one year. Furthermore, capacity releases less than one year can be above the tariff-specified firm toll, while those above a year cannot exceed the firm Tariff rate.
3. We discussed another important feature: recall and reput options. Suppose recall and reput conditions are included in a capacity release transaction. In that case, the original entitlement holder can at any time (within nomination window deadlines) "recall" that capacity and utilize it if they need it. This waives any tolls for the shipper that purchased the capacity release when it was recalled. Once the capacity is recalled, the original

entitlement holder can decide to "reput" that capacity if it no longer needs it and give it back to the shipper that originally purchased the capacity release. Recalls and reputs have to happen on separate days. This is a very common provision.

4. Capacity release above the maximum tariff rate (releases less than or equal to 31 days, greater than or equal to 1 year, or bidded on).
5. The pipeline receives the same reservation charge, whether it be from the releaser or the replacement shipper. However, it could potentially benefit from the volumetric charge if more capacity is used from the capacity release market.

Below, we outline the timelines detailing the processes for conducting open-season auctions for short-term and long-term capacity releases, respectively.

Short Term Release (Less than one year)

1. The capacity release request is posted by 9:00 am on a business day (the specific timing of which will vary depending on the pipeline to meet their nomination windows.
2. The open season/auction period is held between 9:00 and 10:00 am. This is a silent auction; no parties can see what others bid, and bids can be withdrawn during this period.
3. The pipeline company begins evaluation of the bids at 10:00 am, contingencies are cleared, and a determination of the best bid is made based on the process specified by the releaser/requester.
4. Both parties are notified by 11:00 am and confirmed by 12:00 pm. The capacity release contract is awarded within an hour.

Long Term Release (More than one year)

1. The capacity release request is posted by 9:00 am on a business day.
2. The open season/auction period is held for three consecutive business days, such that the open season process ends at 10:00 am three business days after the capacity release posting.
3. The pipeline company begins evaluating the bids at 10:00 am three days later, clearing contingencies and determining the best bid based on the process specified by the releaser/requester.
4. Both parties are notified by 11:00 am and confirmed by 12:00 pm. The capacity release contract is awarded within an hour.

J.2 Interviews With Practitioners

Below, we document our interviews with practitioners who participate in the capacity release market:

- Q: Why would someone not be incentivized to just nominate 100% of their firm contract?

  A: This is due to a certain amount of variable costs associated with transportation. For example, pipeline abandonment surcharges or fuel costs for shipping on the pipeline. If you nominate a particular level of gas for that nomination window, it is assumed you use up to 100% of the gas you nominated. Therefore, you incur variable costs on a volumetric basis equal to your nomination. While these are small in comparison to the overall demand charges (fixed costs), there is little to no incentive for certain types of customers (industrial users or local distribution companies) to incur those costs if they do not require the gas. Whereas marketers/natural gas traders will only choose to nominate the full amount of their contract if the contract path is "in the money": the price of natural gas at the destination plus the variable cost of transportation is greater than the price of natural gas at the origin.

- Q: How often do people choose biddable agreements?

  A: In terms of biddable or non-biddable, he would be surprised if more than 1% of the capacity releases resulted from a biddable process. Most are prearranged deals that you have negotiated before posting them on the pipeline's website. In particular, even in the US, many relationships are established, and many players are aware of who to call if they need a capacity release.

- Q: What is the difference between a natural gas utility, marketer, and retailer?

  A: A natural gas utility is a regulated entity that distributes natural gas, sometimes known as a local distribution company (LDC), to end-use customers such as residential, commercial, and industrial customers. LDCs will often have an exclusive franchise area of a city or region, and in turn, the rates charged for their services are regulated by state regulatory authorities. Utilities source their gas transportation from interstate and intrastate pipelines, which are regulated by federal or state regulatory authorities. These utilities are required by their regulators to hold a certain percentage of their peak day demand requirements in long-term contracts on these pipeline systems. Utilities' average day demand requirements are often much lower than their peak day requirements, and therefore, utilities hold excess transportation capacity on the interstate and intrastate pipeline systems.

Natural gas marketers are unregulated companies that arrange to purchase and sell natural gas. Marketers typically arrange for gas supply agreements with natural gas producers or purchase gas at trading hubs. They then sell this gas to end users or at other market hubs to earn profit based on the spreads between different prices of natural gas hubs. To transport natural gas marketers must also hold capacity on natural gas pipelines but are not required to hold any typical amount of capacity and often choose to be nimble in the amount of capacity they hold on pipeline systems. When arranging the supply of gas for an end user marketers will typically charge for the cost of gas, cost of transportation, and then a service fee for arranging the transactions.

Natural gas retailers are also unregulated companies that arrange to purchase and sell natural gas. However, the main distinction between them and marketers is that retailers typically hold capacity on natural gas distribution systems and sell to end users such as residential, commercial, and industrial customers. The distinction between retailers and LDCs is that the provision of natural gas distribution (the responsibility of LDCs) is regulated, but the retail sales of natural gas are often competitive. End-use consumers can choose to purchase gas directly from the LDC, which offers a regulated rate, or from natural gas retailers, which have competitive rates. Natural gas retailers typically source their natural gas from natural gas marketers or, in certain circumstances, are vertically integrated with a marketing company.

Since marketers and retailers are unregulated entities, they do not require supply obligations and typically hold transportation or distribution capacity that is more in line with the average daily demand for natural gas, in contrast to peak-day demand for utilities.

- Q: Why do we see a large portion of end users of natural gas utilize marketer's services instead of contracting for their own service on interstate natural gas pipelines?

  A: A significant portion of end users of natural gas do not actually hold transportation capacity on interstate pipeline systems, or they are not connected to natural gas distribution utilities. These customers are often industrial facilities, agricultural operations, smaller power generation facilities, and natural gas retailers. Thus, many of these natural gas end users rely on natural gas marketers to meet their gas needs, both in terms of supply and transportation. These end users do not hold transportation capacity on pipelines for two primary reasons: creditworthiness and balance sheet obligations.

  The credit obligations that are necessary for a company to hold pipeline transportation services are often quite steep, and most companies cannot meet them. For example, a typical credit evaluation criterion for firm service on a natural gas pipeline is to provide security guarantees for three months of firm service at the maximum tariff rate for the entire volume of your contract. This requires companies to have large amounts of cash on hand (in the form of an advance deposit), a strong standing letter of credit from a financial institution, an acceptable security interest in collateral, or a guarantee from a more credit-worthy parent company.

As for the balance sheet obligations, given the take-or-pay nature of natural gas transportation firm service contracts, financial institutions view these transportation contracts as debt obligations. If a company were to take out large amounts of transportation capacity, this would result in a large liability appearing on their balance sheet, which may impact their own credit metrics, impacting their ability to secure their own financing and financial obligations.

Given these restrictive requirements, many end users rely on marketers' services to arrange for the supply and transportation of their gas needs since marketers require lower levels of credit requirements and less restrictive outcomes on their balance sheet obligations. This comes at an increased cost to the end user, as marketers often require a service fee or markup for arranging the supply and transportation of natural gas.

- Q: Why do we see natural gas marketers as replacers in the capacity release market?

  A: Since these end users rely on marketers to provide them service when unexpected shocks in demand happen either to end-use residential, commercial, and industrial demand (increased demand for natural gas retailers) or unexpected shocks in their various industries that do not impact demand for gas utilities/retailers (for example an increase in demand for steel production) they often turn to marketers to supply them with additional natural gas. While the primary market for natural gas transportation is held largely by utilities to meet their regulatory obligations, most periods of the year, they do not require the full use of their transportation contracts. They would prefer releasing that capacity to a marketer that will ultimately provide the transportation services to an end user.

K Additional Data Descriptives

K.1 Summary Statistics Regarding IOC and Capacity Release

Here, we show the quarterly proportion of shippers participating in the capacity release market. The data is structured by pipeline, shipper, and quarterly level.

Next, we examine the percentage of shippers who participate in the capacity release market at

Table K.1: Number of Shippers That Are Within the Secondary Market (Quarterly)

Not in the Capacity Release Market	26,388	89%
In the market	3,108	11%
Total	29,496	100%

Table K.2: Number of Shippers that Are Within the Secondary Market (Quarterly, by Pipeline)

Pipeline	**Not in the Capacity Release Market**	**In the Market**	**Proportion**
El Paso	5,854	359	5.78%
Natural Gas	7,479	754	9.16%
Texas	9,681	1,728	15.15%
Transwestern	3,374	267	7.33%
Total	26,388	3,108	10.54%

any point during the contract period.

Table K.3: Number of Shippers That Are Within the Secondary Market

Not in the Capacity Release Market	555	60%
In the market	373	40%
Total	928	100%

Table K.4: Number of Shippers That Are Within the Secondary Market, by Pipeline

Pipeline	**Not in the Capacity Release Market**	**In the Market**	**Proportion**
El Paso	144	66	31.43%
Natural Gas	164	123	42.86%
Texas	144	150	51.02%
Transwestern	103	34	24.82%
Total	555	373	40.19%

Table K.5: Percentage of Contracts With Different Options

Year	Recallable	Reputable	Resale Allowed	Affiliate	Previously Released
2006	72%	58%	72%	0%	13%
2007	75%	63%	81%	0%	5%
2008	79%	67%	85%	0%	8%
2009	94%	72%	93%	0%	9%
2010	94%	66%	92%	0%	10%
2011	93%	65%	92%	1%	12%
2012	94%	70%	94%	1%	12%
2013	95%	68%	94%	0%	9%
2014	95%	61%	94%	0%	7%
2015	95%	65%	94%	0%	7%
2016	97%	78%	97%	0%	6%
2017	96%	81%	97%	0%	6%
2018	96%	80%	98%	0%	5%
2019	97%	80%	97%	0%	6%
2020	97%	79%	96%	0%	10%
2021	98%	78%	96%	0%	18%
2022	98%	79%	96%	0%	20%
2023	99%	80%	96%	0%	20%

Table K.6: List of Pipelines in the Capacity Release Market

Pipeline	Number of Unique Contracts	Number of Unique Releasers	Number of Unique Replacers
ANR Pipeline Company	34,745	164	217
Algonquin Gas Transmission, LLC	93,631	118	240
Colorado Interstate Gas Company, L.L.C.	5,000	84	80
Columbia Gas Transmission, LLC	224,606	344	470
Columbia Gulf Transmission, LLC	47,827	166	303
Crossroads Pipeline Company	4,523	26	71
East Tennessee Natural Gas, LLC	4,740	92	72
Eastern Gas Transmission and Storage, Inc	159,787	241	372
El Paso Natural Gas Company, L.L.C.	5,476	98	161
Enable Gas Transmission, LLC	1,569	156	161
Enable Mississippi River Transmission, LLC	729	16	22
Florida Gas Transmission Company, LLC	26,145	82	111
Gas Transmission Northwest LLC	370	22	34
Great Lakes Gas Transmission Limited Partnership	2,005	39	55
Guardian Pipeline, L.L.C.	359	14	14
Gulf Crossing Pipeline Company LLC	54	3	8
Kern River Gas Transmission Company	2,595	54	86
Midcontinent Express Pipeline LLC	524	18	28
Midwestern Gas Transmission Company	801	36	41
MountainWest Pipeline, LLC	1,483	44	57
Natural Gas Pipeline Company of America LLC	5,616	113	129
Northern Border Pipeline Company	1,510	64	57
Northern Natural Gas Company	15,448	202	190
Northwest Pipeline LLC	4,487	67	80
Panhandle Eastern Pipe Line Company, LP	58,290	116	243
Sea Robin Pipeline Company, LLC	2	1	1
Southern Natural Gas Company, L.L.C.	50,135	264	148
Southern Star Central Gas Pipeline, Inc.	30,154	83	99
Tallgrass Interstate Gas Transmission LLC	85	13	12
Tennessee Gas Pipeline Company, L.L.C.	201,954	334	448
Texas Eastern Transmission, LP	68,795	240	359
Trailblazer Pipeline Company LLC	22	6	8
TransColorado Gas Transmission Company LLC	82	6	7
Transcontinental Gas Pipe Line Company, LLC	159,835	253	350
Transwestern Pipeline Company, LLC	819	31	50
Trunkline Gas Company, LLC	13,428	69	202
Viking Gas Transmission Company	2,665	47	38
Wyoming Interstate Company, L.L.C.	2,657	44	73
Total	1,232,953	3,770	5,097

K.2 Amount Released in Comparison to IOC Data

Table K.7: Amount Released in Comparison to IOC Data (Quarterly)

	p25	p50	p75	Mean	Max
Proportion of quantity released	2.3%	6.8%	25.0%	32.4%	4212.7%

The percentage of the amount released is calculated at the Pipeline-Shipper-Quarterly level, using the Max Daily Transport quantity as a basis.

Table K.8: Amount Released in Comparison to IOC Data (Quarterly)

Pipeline	p25	p50	p75	Mean
El Paso	5.2%	14.2%	25.0%	23.6%
Natural Gas	2.2%	8.2%	25.0%	25.8%
Texas	1.6%	5.3%	13.1%	30.5%
Transwestern	15.2%	25.0%	59.1%	75.4%
Total	2.3%	6.8%	25.0%	32.4%

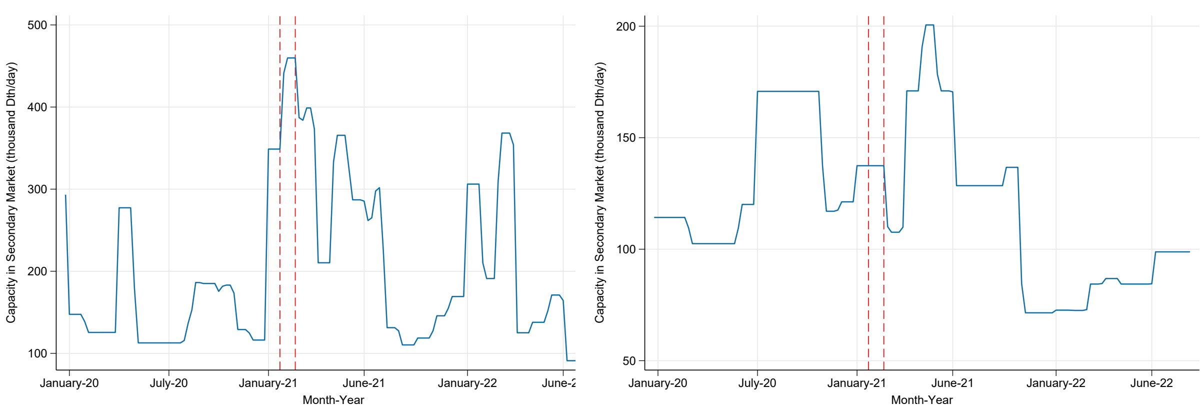

(a) El Paso (b) GTN

Figure K.1: Weekly Capacity Available from Secondary Market

Notes: This figure plots capacity available in the secondary market for each pipeline up to one year before and after the cold wave. Capacity available in the secondary market is defined as capacity contracted from a releasing to a replacement shipper in the secondary market which is available in the current week. The dashed red lines correspond to the official beginning and end of the cold wave.

Curriculum Vitae

Name:	Emmanuel Murray Leclair
Post-Secondary Education and Degrees:	University of Western Ontario London, ON, Canada 2018 - 2024 Ph.D. in Economics
	Carleton University Ottawa, ON, Canada 2017 - 2018 M.A. in Economics
	Université de Montréal Montréal, QC, Canada 2013 - 2016 B.A. in Economics and Political Science
Related Work Experience:	Primary Instructor University of Western Ontario 2022 - 2023
	Graduate Teaching Assistant University of Western Ontario 2018 - 2019
	Graduate Research Assistant University of Western Ontario 2020 - 2024

Fellowships and Awards	Bank of Canada Graduate Student Paper Award (Runner-up) Bank of Canada 2023
	Lasserre-Renzetti Prize for Best Student Paper Canadian Resource and Environmental Economics Association 2023
	SSHRC Productivity Research Fellowship University of Western Ontario 2019-2023
	Western Graduate Research Scholarship University of Western Ontario 2018-2022
	Ontario Graduate Scholarship University of Western Ontario 2020-2022
	Summer Paper Prize University of Western Ontario 2020
	Huntley Schaller Memorial Scholarship in Economics Research Carleton University 2018